



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Simulating Magnetized Laboratory Plasmas with Smoothed Particle Hydrodynamics

J. N. Johnson

July 22, 2009

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Simulating Magnetized Laboratory Plasmas with Smoothed Particle Hydrodynamics

by

JEFFREY N. JOHNSON

B.S. (University of California at San Diego) 1999

M.S. (University of California at Davis) 2004

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Applied Science

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Dr. J. Michael Owen

Professor David Q. Hwang

Professor Garry H. Rodrigue

Chair

Committee in Charge

2009

Simulating Magnetized Laboratory Plasmas with Smoothed Particle Hydrodynamics

Abstract

The creation of plasmas in the laboratory continues to generate excitement in the physics community. Despite the best efforts of the intrepid plasma diagnostics community, the dynamics of these plasmas remains a difficult challenge to both the theorist and the experimentalist. This dissertation describes the simulation of strongly magnetized laboratory plasmas with Smoothed Particle Hydrodynamics (SPH), a method born of astrophysics but gaining broad support in the engineering community. We describe the mathematical formulation that best characterizes a strongly magnetized plasma under our circumstances of interest, and we review the SPH method and its application to astrophysical plasmas based on research by Phillips [1], Børve [2], and Price and Monaghan [3]. Some modifications and extensions to this method are necessary to simulate terrestrial plasmas, such as a treatment of magnetic diffusion based on work by Brookshaw [4] and by Atluri [5]; we describe these changes as we turn our attention toward laboratory experiments. Test problems that verify the method are provided throughout the discussion. Finally, we apply our method to the compression of a magnetized plasma performed by the Compact Toroid Injection eXperiment (CTIX) [6] and show that the experimental results support our computed predictions.

To Christina, for putting up with my nonsense.

Contents

List of Figures	vii
List of Tables	xii
1 Introduction	1
1.1 Plasmas and Magnetic Fields	1
1.1.1 Large-Scale Magnetic Fields in Space	1
1.1.2 Laboratory Plasmas	2
1.2 The Simulation of Magnetized Plasmas	3
1.2.1 Numerical Difficulties	3
1.2.2 A Brief History	5
1.3 Mesh-Free Methods	6
1.3.1 Smoothed Particle Hydrodynamics	7
1.3.2 Mesh-Free Galerkin and Galerkin-Like Methods	7
1.3.3 Least-Squares Approaches	9
1.4 Features of the Method	10
1.4.1 Limitations	12
1.5 Dissertation Roadmap	13
2 The Physics of Magnetized Plasmas	14
2.1 Charged Particles in Magnetic Fields	14
2.1.1 Plasma oscillations	15
2.2 Mathematical description	16
2.2.1 Plasma Kinetic Theory	18
2.2.2 Fluid Equations	19
2.3 The Magnetohydrodynamic Approximation	21
2.3.1 The Induction Equation	22
2.3.2 The Single-Fluid Equations	24
2.3.3 Boundary Conditions at Material Interfaces	25
2.3.4 Lagrangian Formulation	26
2.4 The Physics of MHD	27
2.4.1 MHD Waves	27
2.4.2 The Behavior of the Magnetic Field	29
2.5 Limitations of MHD	31
2.6 Requirements for Modeling Plasmoids	32

3	Smoothed Particle Hydrodynamics	35
3.1	The SPH Formulation	35
3.1.1	Spatial discretization	39
3.1.2	The fluid equations	43
3.1.3	The SPH smoothing scale	44
3.2	Shock Waves	46
3.2.1	The Rankine-Hugoniot Relations	47
3.2.2	Artificial Viscosity	50
3.3	Time Integration	53
3.3.1	Testing Convergence	55
3.3.2	Constraints on Δt	56
3.4	Issues and Refinements	58
3.4.1	Node Disorder	58
3.4.2	Stability	59
3.4.3	Improved energy conservation	60
3.5	Multidimensional SPH	66
3.5.1	Artificial Viscosity	66
3.5.2	Computational Expense	67
3.6	Boundary Conditions	67
3.6.1	Reflecting and rigid boundaries	68
3.6.2	Periodic boundaries	68
3.6.3	Inflow and outflow boundaries	69
3.6.4	Dirichlet boundaries	70
3.6.5	Neumann boundaries	70
3.6.6	Free boundaries	71
3.6.7	Cylindrically symmetric systems	72
3.7	Curved Surfaces	73
3.7.1	Rigid Body Treatment	74
3.7.2	The Single and Multiple Tangent Boundary Methods	75
3.7.3	Simple Reflection and Scaling Method	76
3.8	Test Problems	77
3.8.1	Hydrostatic equilibrium	78
3.8.2	An acoustic wave	80
3.8.3	The Sod shock tube	86
3.8.4	The Noh problem	101
3.9	Summary	117
4	Magnetic Fields	118
4.1	Smoothed Particle Ideal Magnetohydrodynamics	118
4.1.1	The $\nabla \cdot \mathbf{B} = 0$ constraint	120
4.1.2	The Tensile Instability	122
4.1.3	Smoothing field discontinuities	123
4.1.4	The Compatible Energy Method Extended to MHD	124
4.1.5	Boundary Conditions on \mathbf{B}	125
4.2	Magnetic Diffusion in SPH	126
4.2.1	Laplacian Operators in SPH	127
4.2.2	Boundary Conditions	132
4.2.3	Limitations	134
4.3	A Meshless Petrov-Galerkin Approach to Magnetic Diffusion	134
4.3.1	The Local Weak Form	135
4.3.2	Semi-Discrete Equations	136

4.3.3	Meshless Shape Functions	138
4.3.4	Fully Discrete Equations	141
4.3.5	Boundary Conditions	142
4.3.6	Limitations	143
4.4	Resistive Smoothed Particle Magnetohydrodynamics	144
4.4.1	Stencil Coupling	145
4.4.2	Operator Splitting	146
4.5	Test Problems	146
4.5.1	The decay of a magnetic pulse	147
4.5.2	The magnetic “wine cellar” problem	152
4.5.3	Continuously-varying resistivity	153
4.5.4	Discontinuous resistivity	155
4.5.5	Time-dependent resistivity	157
4.5.6	Decay of a 2D magnetic flux tube	158
4.5.7	The advection of a magnetic pulse	159
4.5.8	Magnetosonic waves	164
4.5.9	Alfvén waves	170
4.5.10	Brio and Wu’s shock tube	172
4.5.11	Planar magnetized compression	180
4.5.12	Magnetic rarefaction	185
4.5.13	Spherical Blast Wave	189
4.5.14	MHD Turbulence: the Orzsag-Tang vortex	194
4.6	Summary	198
5	Simulations	200
5.1	The Device	200
5.1.1	Diagnostics	202
5.2	Plasma Compression	205
5.3	Simulating the Compression of a Plasmoid	206
5.3.1	Initial Conditions: Force-Free Equilibrium	207
5.3.2	The Plasma/Vacuum Interface	209
5.3.3	Plasma Resistivity	210
5.3.4	Artificial Dissipation and Spurious Viscous Heating	210
5.3.5	Results	211
5.4	Conclusion	223
6	Conclusion	224
6.1	A Critical Assessment of SPMHD	224
6.2	Extensions	226
6.3	Further Study of CTIX	226
6.4	Directions for Improvement	227
	Bibliography	230
A	Specific Forms of MLPG Matrices	237
B	Software Environment	240

List of Figures

2.1	The waves present in (a) ρ , (b) v_x , (c) v_y , (d) u , (e) p , and (f) B_y in an MHD shock tube: (i) left-moving fast rarefaction, (ii) left-moving slow compound wave (shock/rarefaction), (iii) contact discontinuity, (iv) right-moving slow shock, and (v) right-moving fast shock.	30
3.1	The “stencil” of a node in SPH in 2D. The radius of influence is determined by the desired number of neighbors per smoothing scale, n_h : (a) $n_h = 1.2$, (b) $n_h = 1.4$, (c) $n_h = 1.8$, (d) $n_h = 2$	36
3.2	The gradual steepening of a linear wave. (a) is the original acoustic wave, and (b) and (c) show the distortion of the wave through the $\mathbf{v} \cdot \nabla \mathbf{v}$ term.	46
3.3	The behavior of (a) h , the SPH smoothing scale, and (b) Ω , the ∇h correction factor, near a shock.	62
3.4	Ghost nodes for a set of interior nodes for (a) rigid and (b) periodic boundaries.	68
3.5	The positions of nodes initially in a cubic lattice after thermal motion with $n_h = 1.8$. The interior nodes seem to have settled into a close-packed lattice.	80
3.6	An acoustic wave after having passed through the domain (a) once ($t = 1$) and (b) five times ($t = 5$), using the unsymmetric energy equation and $n_h = 2$	82
3.7	The convergence rates of (a) the mass density and (b) the velocity of an acoustic wave having crossed a domain five times using the unsymmetric energy equation and $n_h = 2$	83
3.8	An acoustic wave <i>without artificial viscosity</i> after having passed through the domain (a) once ($t = 1$) and (b) five times ($t = 5$), using the unsymmetric energy equation and $n_h = 2$	84
3.9	An acoustic wave at (a) ($t = 1$) and (b) ($t = 5$) using the unsymmetric energy equation and $n_h = 2$ with tempered artificial viscosity coefficients $\alpha = 1/4$ and $\beta = 0$	85
3.10	The convergence rates of (a) the mass density and (b) the velocity of the acoustic wave at $t = 5$ using the unsymmetric energy equation and $n_h = 2$ with tempered artificial viscosity coefficients $\alpha = 1/4$ and $\beta = 0$	85
3.11	The initial conditions for the Sod shock tube problem: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	87
3.12	The solution of the Sod shock tube problem at $t = 0.15$ using the unsymmetric energy equation: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	88
3.13	Convergence rates for the Sod problem using the unsymmetric energy equation with $n_h = 2$: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	89
3.14	Convergence rates for the Sod problem using the unsymmetric energy equation with $n_h = 2$ and ∇h corrections: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	91
3.15	Convergence rates for the Sod problem using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	92
3.16	The specific energy profile for the Sod problem at $t = 0.15$ using (a) ∇h corrections and (b) the compatible energy method.	93

3.17	The ∇h correction factor Ω for the Sod problem at $t = 0.15$	94
3.18	Convergence rates for the Sod problem using the compatible energy method with $n_h = 2$ and with ∇h corrections: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy . .	95
3.19	The solution (a) and convergence (b) of the specific thermal energy in the Sod problem at $t = 0.15$ using the compatible energy method and $n_h = 2$ with Price's artificial heat conduction terms.	96
3.20	The solution of the Sod shock tube problem at $t = 0.15$ using the compatible energy method and $n_h = 2$ with variable artificial viscosity: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	97
3.21	The convergence of the Sod shock tube problem at $t = 0.15$ using the compatible energy method and $n_h = 2$ with variable artificial viscosity: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	98
3.22	The solution of the Sod shock tube problem at $t = 0.15$ using the compatible energy method and $n_h = 2$, compared with that computed by Athena: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	99
3.23	The convergence of the Sod shock tube problem at $t = 0.15$ using the compatible energy method and $n_h = 2$, using Athena's solution as the accepted answer: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	100
3.24	Solution of the Noh problem at $t = 0.3$ using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	102
3.25	Convergence rates for the planar Noh problem using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy	103
3.26	Solutions to the cylindrical Noh problem at $t = 0.6$ using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy	106
3.27	Convergence rates for the cylindrical Noh problem using the unsymmetric energy equation with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy	107
3.28	Plots of nodes in the cylindrical Noh problem at $t = 0.6$ using the unsymmetric energy equation with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy	108
3.29	Solutions to the cylindrical Noh problem at $t = 0.6$ using the compatible energy method and ∇h corrections with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy	110
3.30	Plots of nodes in the spherical Noh problem at $t = 0.6$ using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_r	111
3.31	Convergence rates for the spherical Noh problem using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy	112
3.32	(a) Initial velocity configuration of nodes in the conical Noh problem, and (b) leakage of nodes through the origin in the absence of a rigid plane at $z = 0$	114
3.33	Plots of nodes in the conical Noh problem at $t = 0.6$ using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_r	114
3.34	Convergence rates for the conical Noh problem using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy	116
4.1	Initial configuration for (a) B and (b) J in the decay of a square magnetic pulse.	148
4.2	Computed and analytic solutions for (a) B and (b) J in the decay of a square magnetic pulse at $t = 0.5$	149
4.3	Convergence of solutions for (a) B and (b) J in the decay of a square magnetic pulse at $t = 0.5$	149
4.4	Computed and analytic solutions for (a) B and (b) J in the decay of a Gaussian magnetic pulse at $t = 0.5$	150
4.5	Convergence of solutions for (a) B and (b) J in the decay of a Gaussian magnetic pulse at $t = 0.5$	151

4.6	(a) Initial ($t = 0$) and (b) final ($t = 0.5$) solutions for $ \mathbf{B} $ in the two-dimensional decay of a square magnetic pulse.	151
4.7	Computed and analytic solutions for (a) \mathbf{B} and (b) \mathbf{J} in the magnetic “wine cellar” problem at $t = 0.1$	152
4.8	Convergence of solutions for (a) \mathbf{B} and (b) \mathbf{J} in the magnetic “wine cellar” problem at $t = 0.1$	153
4.9	Convergence of solutions for (a) \mathbf{B} and (b) \mathbf{J} in the magnetic “wine cellar” problem with perturbed nodes at $t = 0.1$	154
4.10	(a) Solution for \mathbf{B} in a conductor with continuously-varying diffusivity at $t = 0.1$ and (b) the convergence of its L_p error norms for $N = 100, N = 200, N = 400, N = 800$	155
4.11	(a) Solution for \mathbf{B} in a region with discontinuous diffusivity at $t = 1$ using (4.26), and (b) the convergence of its L_p error norms for $N = 100, N = 200, N = 400, N = 800$	156
4.12	(a) Solution for \mathbf{B} in a region with discontinuous diffusivity at $t = 1$ using the product rule and (3.23) for $\nabla \cdot (D\nabla\mathbf{B})$, and (b) the convergence of its L_p error norms for $N = 100, N = 200, N = 400, N = 800$	157
4.13	(a) Solution for \mathbf{B} for time-varying diffusivity $D(t) = D_0 e^{-t/\tau}$ at $t = 0.01$ and (b) the convergence of its L_p error norms for $N = 100, N = 200, N = 400, N = 800$	158
4.14	Initial and final configurations for \mathbf{B} and \mathbf{J} in the cylindrical magnetic flux tube problem. The artifacts along the x and y axes appear both in the initial and evolved solutions and result from interpolation to a structured mesh. (a) $\mathbf{B}(x, 0)$, (b) $J_z(x, 0)$, (c) $\mathbf{B}(x, 2)$, (d) $J_z(x, 2)$	160
4.15	(a) Comparison of computed B_θ (dots) with exact solution (ticks) for the cylindrical magnetic flux tube problem. (b) The convergence of $L_p(\mathbf{B})$ for this problem.	161
4.16	Advection of a square magnetic pulse in a perfect conductor with $v = 1$ (a) initially, (b) halfway across the domain, and (c) after a crossing. Since SPH is Lagrangian, there is no degradation in the signal.	161
4.17	Current density of the advecting magnetic pulse (a) initially, (b) halfway across the domain, and (c) after a crossing.	162
4.18	Advection of a square magnetic pulse in a perfect conductor with $v = 1$ (a) initially, (b) halfway across the domain, and (c) after a crossing. The decay in the pulse is caused only by the magnetic diffusion equation.	162
4.19	Current density of the advecting magnetic pulse (a) initially, (b) halfway across the domain, and (c) after a crossing.	163
4.20	Convergence of the solution for (a) \mathbf{B} and (b) \mathbf{J} in the advection of a magnetic pulse through a medium with diffusivity $D = 0.01$	163
4.21	Density profile for fast magnetosonic wave after (a) 1 crossing and (b) 5 crossings with $n_h = 2$ using the unsymmetric energy equation ($\alpha = \alpha_B = 0.1$)	165
4.22	Convergence of solutions for (a) ρ (b) \mathbf{v} , and (c) \mathbf{B} for fast magnetosonic wave after 5 crossings with $n_h = 2$ using the unsymmetric energy equation ($\alpha = \alpha_B = 0.1$)	166
4.23	Convergence of solutions for (a) ρ , (b) \mathbf{v} , and (c) \mathbf{B} for fast magnetosonic wave after 5 crossings with $n_h = 2$ using the ∇h corrections to the unsymmetric energy equation ($\alpha = \alpha_B = 0.1$)	168
4.24	Convergence of solutions for (a) ρ (b) \mathbf{v} and (c) \mathbf{B} for fast magnetosonic wave after 5 crossings with $n_h = 2$ using the compatible energy equation ($\alpha = \alpha_B = 0.1$)	169
4.25	Solution for B_y (a) after 1 crossing and (b) after 5 crossings for circularly-polarized Alfvén wave with $n_h = 2$ using the unsymmetric energy equation	171
4.26	Convergence of solutions for (a) \mathbf{v} and (b) \mathbf{B} for circularly-polarized Alfvén wave after 5 crossings with $n_h = 2$ using the unsymmetric energy equation	171
4.27	Alfvén wave after 5 crossings with (a) $n_h = 1.2$, (b) $n_h = 1.4$, (c) $n_h = 1.6$, and (d) $n_h = 1.8$ using the uncorrected unsymmetric energy equation	173
4.28	Onset of the tensile instability at $t = 0.1$ in the mass density: (a) $n_h = 1.6$, (b) $n_h = 1.8$. Note the clumping of the nodes around the contact discontinuity.	175
4.29	Onset of the tensile instability at $t = 0.1$ in the velocity: (a) $n_h = 1.6$, (c) $n_h = 1.8$. Slight oscillations start to appear to the right of the contact discontinuity.	176

4.30	Solutions to the Brio/Wu shock tube problem ($n_h = 1.6$) at $t = 0.1$ using the unsymmetric energy equation: (a) ρ , (b) v_x , (c) v_y , (d) u , (e) p , (f) B_y .	177
4.31	Solutions to the Brio/Wu shock tube problem ($n_h = 1.6$) at $t = 0.1$ using the compatible energy method and artificial thermal conduction with a variable α_u : (a) ρ , (b) v_x , (c) v_y , (d) u , (e) p , (f) B_y .	178
4.32	Solutions to the Brio/Wu shock tube problem ($n_h = 1.6$) at $t = 0.1$ using the compatible energy method, artificial thermal conduction (variable α_u), and ∇h corrections: (a) ρ , (b) v_x , (c) v_y , (d) u , (e) p , (f) B_y .	179
4.33	Solutions to the magnetic compression problem ($n_h = 2$) at $t = 0.08$ using the unsymmetric energy equation: (a) ρ , (b) v_x , (c) v_z , (d) u , (e) p , (f) B_z .	182
4.34	Convergence rates for the L_p norms of the solutions to the magnetic compression problem ($n_h = 2$) using the unsymmetric energy equation: (a) ρ , (b) \mathbf{v} , (c) u , (d) \mathbf{B} .	183
4.35	Solution for v_z in the planar magnetic compression problem using (a) the compatible energy method and (b) the ∇h corrections.	184
4.36	Solution for B_z in the planar magnetic compression problem using (a) the compatible energy method and (b) the ∇h corrections.	184
4.37	Solutions to the magnetic rarefaction problem ($n_h = 2$) at $t = 0.1$ using the unsymmetric energy equation: (a) ρ , (b) v_x , (c) u , (d) p .	187
4.38	Convergence rates for the L_p norms of the solutions to the magnetic compression problem ($n_h = 2$) using the unsymmetric energy equation: (a) ρ , (b) \mathbf{v} , (c) u , (d) p .	188
4.39	Solution for ρ in the magnetic rarefaction problem using (a) the compatible energy method and (b) the ∇h corrections.	189
4.40	Solutions to the Balsara blast wave at $t = 0.02$: (a) ρ (Athena), (b) ρ (SPH), (c) $ \mathbf{B} $ (Athena), (d) $ \mathbf{B} $ (SPH).	191
4.41	Convergence rates for the L_p norms of the Balsara blast wave at $t = 0.02$ using the ∇h corrections and the compatible energy method: (a) ρ , (b) \mathbf{v} , (c) u , (d) \mathbf{B} .	192
4.42	Self-convergence rates for the L_p norms of Athena's solutions to the Balsara blast wave at $t = 0.02$: (a) ρ , (b) \mathbf{v} , (c) p , (d) \mathbf{B} .	193
4.43	Solutions to the Orszag-Tang vortex at $t = 0.5$: (a) ρ (Athena), (b) ρ (SPH), (c) $ \mathbf{B} $ (Athena), (d) $ \mathbf{B} $ (SPH).	195
4.44	Self-convergence rates for the L_p norms of Athena's solutions to the Orszag-Tang vortex at $t = 0.5$: (a) ρ , (b) \mathbf{v} , (c) p , (d) \mathbf{B} .	196
4.45	Convergence rates for the L_p norms of the Orszag-Tang vortex at $t = 0.5$ using the uncorrected energy equation: (a) ρ , (b) \mathbf{v} , (c) u , (d) \mathbf{B} .	197
5.1	The CTIX device (compression nozzle not shown)	201
5.2	The toroidal pushing field, distinguished from the plasma's internal field. The pushing field is created by electric current, which flows through the outer electrode, passes through the conducting plasma, and returns along the inner electrode.	202
5.3	The compression nozzle with ports for magnetic probes.	203
5.4	(a) B_z and (b) B_z in a well-formed compact toroid as it passes probes at different axial points in the acceleration section. Taken from CTIX shot 71310.	204
5.5	The initial configuration of the plasma as rendered by the Visit visualization tool. The nozzle boundary is represented by a purely cosmetic uniform structured mesh.	212
5.6	The development of a large magnetic field along the axis pushes the plasma outward along the nozzle. (a) shows the magnetic field and (b) the plasma's mass density.	213
5.7	Experimental photographs of the nozzle after the plasma compressions. Shots (a) 71310, (b) 71336, (c) 71337, and (d) 71361.	215
5.8	The final configuration of a plasma with $B_0 = 5000$ G that does not escape from the nozzle. (a) mass density, and (b) magnetic field intensity.	216

5.9	The emergence of a plasma with $B_0 = 2000$ G from the nozzle. The magnetic field growth along the axis is still present, but the kinetic energy of the plasma allows it to overcome the magnetic pressure. (a) mass density, and (b) magnetic field intensity.	217
5.10	Schematic representation of <i>poloidal</i> magnetic reconnection in a compact toroid that is expanding radially and moving to the right. The poloidal (axial) field lines are aligned as they approach one another, reconnect, and produce a more intense field.	219
5.11	Magnetic field amplitude in 2D simulation of a cross section of an expanding CT. (a) initial configuration, and (b) after reconnection ($t = 0.1\mu s$).	219
5.12	Schematic representation of toroidal magnetic reconnection. The toroidal field lines are opposed as they approach one another, annihilating one another upon reconnection.	220
5.13	Magnetic field amplitude in 2D simulation of a cross section of an expanding jet. (a) initial configuration, and (b) after reconnection ($t = 0.1\mu s$).	221
5.14	(a) B_z and (b) B_θ in a toroidal plasma jet as it passes probes at different axial points in the acceleration section. This is experimental data taken from CTIX shot 71356. Contrast these signals with those in Figure 5.4.	221
5.15	Experimental photographs of plasma jets emerging from the nozzle at right. Shots (a) 71356, (b) 71359, (c) 71360, and (d) 71375.	222

List of Tables

3.1	Global error $ \tilde{y} - y(1) $ for FE (Forward Euler), RK2 (second-order Runge-Kutta), PC (Predictor-Corrector), LF (Leapfrog), and RK4 (fourth-order Runge-Kutta) integrators as a function of number of steps N	56
3.2	Convergence rates for global error $ \tilde{y} - y(1) $ for FE (Forward Euler), RK2 (second-order Runge-Kutta), PC (Predictor-Corrector), LF (Leapfrog), and RK4 (fourth-order Runge-Kutta) integrators	56
3.3	$L_1(\rho)$ in 1D hydrostatic equilibrium for various n_h	79
3.4	Convergence rates for L_p norms of the mass density in the acoustic wave as functions of n_h	83
3.5	Convergence rates for L_p norms of the velocity in the acoustic wave as functions of n_h	83
3.6	Convergence rates for L_p norms of the mass density in the Sod problem as functions of n_h	90
3.7	Convergence rates for L_p norms of the velocity in the Sod problem as functions of n_h	90
3.8	Convergence rates for L_p norms of the pressure in the Sod problem as functions of n_h	90
3.9	Convergence rates for L_p norms of the specific thermal energy in the Sod problem as functions of n_h	90
3.10	Convergence rates for L_p norms of the mass density in the planar Noh problem using the various energy schemes.	104
3.11	Convergence rates for L_p norms of the velocity in the planar Noh problem using the various energy schemes.	104
3.12	Convergence rates for L_p norms of the specific thermal energy in the planar Noh problem using the various energy schemes.	104
3.13	Convergence rates for L_p norms of the pressure in the planar Noh problem using the various energy schemes.	105
3.14	Convergence rates for L_p norms of the mass density in the cylindrical Noh problem using the various energy schemes.	105
3.15	Convergence rates for L_p norms of the velocity in the cylindrical Noh problem using the various energy schemes.	106
3.16	Convergence rates for L_p norms of the specific thermal energy in the cylindrical Noh problem using the various energy schemes.	107
3.17	Convergence rates for L_p norms of the pressure in the cylindrical Noh problem using the various energy schemes.	109
3.18	Convergence rates for L_p norms of the mass density in the spherical Noh problem using the various energy schemes.	110
3.19	Convergence rates for L_p norms of the velocity in the spherical Noh problem using the various energy schemes.	111
3.20	Convergence rates for L_p norms of the specific thermal energy in the spherical Noh problem using the various energy schemes.	111

3.21	Convergence rates for L_p norms of the pressure in the spherical Noh problem using the various energy schemes.	112
3.22	Convergence rates for L_p norms of the mass density in the conical Noh problem using the various energy schemes.	115
3.23	Convergence rates for L_p norms of the velocity in the conical Noh problem using the various energy schemes.	115
3.24	Convergence rates for L_p norms of the specific thermal energy in the conical Noh problem using the various energy schemes.	115
3.25	Convergence rates for L_p norms of the pressure in the conical Noh problem using the various energy schemes.	115
4.1	Convergence rates for L_p norms of the mass density in the magnetosonic wave as functions of n_h	167
4.2	Convergence rates for L_p norms of the velocity in the fast magnetosonic wave as functions of n_h	167
4.3	Convergence rates for L_p norms of the magnetic field in the fast magnetosonic wave as functions of n_h	167
4.4	Convergence rates for L_p norms of \mathbf{v} in the 1D Alfvén wave as functions of n_h	171
4.5	Convergence rates for L_p norms of \mathbf{B} in the 1D Alfvén wave as functions of n_h	172
4.6	Convergence rates for L_1 error norms of the solution for the Brio/Wu shock tube as a function of n_h . The unsymmetric energy equation was used to produce these results.	175
4.7	Convergence rates for L_1 error norms of the solution for the Brio/Wu shock tube as a function of n_h . These results were produced using the compatible energy method and artificial thermal conduction with Morris's variable switch α_u	180
4.8	Convergence rates for L_1 error norms of the solution for the Brio/Wu shock tube as a function of n_h . These results were produced using the compatible energy method, artificial thermal conduction (with a variable α_u), and ∇h corrections.	180
4.9	Convergence rates for L_1 error norms of the solution for the magnetic compression problem using the different methods for the energy equations.	185
4.10	Convergence rates for L_1 error norms of the solution for the magnetic rarefaction problem using the different methods for the energy equations.	187
4.11	Convergence rates for L_1 error norms of the solution for Balsara's blast wave using the different methods for the energy equations.	190
4.12	Convergence rates for L_1 error norms of the solution for the Orszag-Tang vortex using the different methods for the energy equations.	196
5.1	Estimated peak magnetic field measurements for shots in Figure 5.7. These fields are measured as the plasma enters the nozzle.	214
5.2	Results of simulations with various axial velocities and internal magnetic field strengths. The "pass" column indicates whether the plasma passes successfully through the nozzle.	218
5.3	Estimated peak magnetic field measurements for shots in Figure 5.15. These fields are measured as the plasma enters the nozzle.	223

Acknowledgements

A number of people remained my friends and helped me through this dissertation despite my insufferable moments. First and foremost, my committee: a big thank-you to Mike Owen, my technical advisor, who took me on as a student and walked me through the process of studying hydrodynamics and whose guidance and discipline have been indispensable during my trek through that most nebulous nexus of SPH and magnetohydrodynamics. I am indebted to David Hwang for the data obtained from his experiment against which I validated my work, and for his support in working in Hertz hall alongside the group. To Garry Rodrigue, my thesis advisor: thank you for your help in understanding numerical mathematics, and for your support during various (and varied) phases of my matriculation at UC Davis.

I would like to thank Doug Miller, who selflessly read through scads and scads of code, and myriad posters and text on my behalf, for countless pieces of good advice, and invariably good reading recommendations. Thanks to Nick Gentile, who gave me a copy of Bill Noh's paper before I knew what I was getting into, and for many rockin' evenings throughout the last several years. To Mike Zika and Lila Chase, for their encouragement and support during my tenure as an LLNL employee, and to Paul Renard for countless thoughtful conversations and observations: thank you.

A big thanks also to the CTIX crew, present and past: to Robert Horton, whose physical intuition, experimental acumen, and intellectual tenacity are an inspiration; to Russ Evans, for all of his help with data acquisition and for some much-needed perspective (and break-time); to Glenn Umont, for his constant encouragement and relaxed disposition in the face of science; to Stephen Howard, for the use of his extremely thorough thesis and for his guidance early on.

I also owe the riders of the 9:09 Wheels 20X a heartfelt thanks for making the commute to Livermore from San Francisco a good deal more enjoyable than I thought possible: José, Aleks, Miriam, Brenda, Alastair, Greg, Jonathon, and all those fairweather riders that pop in from time to time. Also: to the now-gainfully-employed Andrew—can you please find José a job at Berkeley?

There are a few (more) DAS students to whom I extend my thanks: to Rob Rieben, for helping

me learn about the Vector Finite Element method and for supplying his thesis as a template, and for hanging around in Oakland all those years back. Thanks also to Aaron Fisher, one of the other “last ones out,” for his advice, for good times, and for the handsome acid-free thesis paper!

I owe a debt of gratitude to Professor Niels Jensen and to Tiffany Macias-Silva at DAS for their encouragement, their understanding, and for helping me navigate the process of filing this thesis.

Thanks go to my horn teacher Steven Keys for taking me seriously as a musician and for giving me a much-needed distraction from physics this last year. Thanks also to the members of Brassavola, Zaptet, Octangle, the Rainbow Chamber Players, and the Bay Area Rainbow Symphony for giving me a good excuse to pursue this distraction!

In particular, I would like to thank my family, for all of their support throughout the years. Thanks to my mother, for providing me with the tenacity to finish the bloody thing; to my father, for getting me interested in science in the first place; to my brother Scott, for his encouragement and for some great virtual venting of very real frustrations.

Finally, to my lovely wife: Thank you, Christina, for your patience and flexibility in sharing me with this thesis.

Chapter 1

Introduction

1.1 Plasmas and Magnetic Fields

Classical physics has long been understood in principle. Newton's Laws of motion were written down in 1687 [7] and Maxwell expressed his theory of electromagnetism in 1873 [8] in a way that has remained largely intact since. To this day these theories continue to be used to study systems of increasing complexity. Plasma physics, the study of ionized particles in electromagnetic fields, is a perfect example of the complexity that can be achieved by combining Newton's force laws with Maxwell's equations to study the dynamics of systems with astronomically large numbers of particles ($N \sim 10^{15}$ and upward). In these large systems, it is simply not possible to understand the interactions of particles in simple analytic terms. The fact that over 99 percent of the observable baryonic material in the universe is plasma provides ample motivation for modeling these interactions through theory, experiment, and computation.

1.1.1 Large-Scale Magnetic Fields in Space

Magnetic fields have become an item of interest in astrophysics and cosmology, since they are thought to play important roles in several phenomena. These fields contribute to the formation of planets and stars and to the dynamics of star mergers by channeling the flow of molecular clouds. They constrain the

motion of charged matter in stellar phenomena like solar flares [9]. They have been found in the vicinity of white dwarfs [10] and neutron stars [11]. Magnetic dynamos, which are magnetized configurations whose fields are amplified by turbulent fluid motion, are thought to contribute to instabilities and outflows in accretion discs [12]. The Earth itself has a dynamic magnetic field whose origin is still not well understood but may be explained in part by such a dynamo [9].

Such developments have motivated further interest in creating laboratory experiments to study strong magnetic fields in plasmas ([13], [14]), but the field of laboratory astrophysics itself is hardly new. It dates back to Birkeland’s work in understanding the aurora borealis in the early twentieth century [15], and generated some excitement in the mid-century with Laurence’s plasma “test-tube universe” [16]. For decades since, laboratory tools have been used to create structures analogous to those observed in space: for example, pinches [17] and plasma focuses [18].

1.1.2 Laboratory Plasmas

In addition to their role in studying astrophysics, plasmas formed in the laboratory have several terrestrial applications. Magnetized laboratory plasmas are most notably associated with magnetic fusion ([19]). The magnetic fusion field itself has driven several areas of research, including plasma formation through dynamos [20], confinement and acceleration by magnetic fields [21], hydromagnetic stability, and the development of plasma diagnostics [22].

Central to each of these topics is the notion of a plasma which can be controlled. Since charged particles can be made to orbit magnetic field lines, it makes sense for a plasma to have an internal magnetic field, reducing the rate of plasma expansion. Thus, the idea of a self-contained magnetized plasma, or *plasmoid*, is attractive. Plasmoids were first created by Bostick ([23], [24]) and were used to study terrestrial phenomena like ball lightning, as well as astrophysical phenomena like the solar wind. Plasmoids exist with varying magnetic field intensities, and in various configurations. In this thesis we confine our interest to a class called *compact toroids* (CTs), which are strongly magnetized rings of plasmas formed within a conducting vessel of coaxial geometry [25]. A CT has currents that sustain its internal field, and this internal field is strong enough

that the plasma can be accelerated by an external current source. We reserve a detailed discussion of compact toroids for Chapter 5.

1.2 The Simulation of Magnetized Plasmas

As we have discussed, plasma experiments have already helped us to understand several natural phenomena. Thus, there is great interest in the design of such experiments in the plasma community. However, it remains a difficult problem to quantify and predict the behavior of laboratory plasmas. While the development of plasma diagnostics has seen many advances in recent years [22], many experiments still have problems precisely reproducing and measuring basic quantities in plasmas.

Numerical modeling is an important supplementary tool in experimental design. It connects plasma theory with non-idealized systems so that our intuition of these systems may be improved. In this thesis we describe a numerical method for modeling plasmas in the U. C. Davis Compact Toroid Injection eXperiment [6], an experiment that has been used to study plasma acceleration, magnetic reconnection, MHD turbulence, and the development of several diagnostics.

1.2.1 Numerical Difficulties

The modeling of a magnetized plasma poses a serious challenge to any numeric technique. To begin with, the equations of gas dynamics are difficult to solve because the presence of shocks and contact discontinuities require special treatment. These features render useless many techniques that are fruitful in treating other physical systems. For instance, high-order methods used to resolve linear acoustic waves are reduced to lower orders because the derivatives they use to increase their accuracy simply do not exist at a discontinuity. Further, one must describe not only the motion of the plasma itself, but also the evolution of the magnetic field.

Magnetic fields have several features that make them difficult to represent using conventional finite difference and finite element methods. Firstly, they are inherently three-dimensional, even in systems which

vary in only one or two coordinate dimensions. This is particularly evident in terms like the $\nabla \times (\mathbf{v} \times \mathbf{B})$ term that appears in treatments of moving conductors in magnetic fields, and in the $\mathbf{J} \times \mathbf{B}$ Lorentz force. The components of velocity transverse to the flow do not appear in the equations of ordinary fluid mechanics.

Secondly, magnetic fields are manifestly divergence-free. Gauss’s Law of electrostatics relates the divergence of an electric field to the presence of electric charge in space:

$$\nabla \cdot \mathbf{E} = \frac{\rho_e}{4\pi\epsilon_0}. \quad (1.1)$$

Gauss’s Law of Magnetism does the same for magnetic fields. However, no isolated magnetic charges have been discovered in nature, so

$$\nabla \cdot \mathbf{B} = 0. \quad (1.2)$$

That is, magnetic field lines close on each other: there are no “sources” of field lines. Any treatment of magnetic fields cannot justifiably ignore this topological property—the results of doing so can be catastrophic [26].

Finally, magnetic fields obey evolution equations that are not strictly hyperbolic. Hyperbolic systems of differential equations (e.g. wave/advection equations) involve the propagation of a signal at some finite speed, in which parts of a domain are causally-connected and disconnected. These propagation speeds are important for certain numerical methods like Godunov methods [27], since they determine the domains of a signal’s influence. A *strictly hyperbolic* system of equations has a Jacobian matrix with real and distinct eigenvalues (corresponding to distinct signal speeds). MHD complicates this structure, since it has waves at several propagation speeds—the speeds of these waves can be the same in some directions. This degeneracy makes the system more difficult to analyze. For example, it is only relatively recently that Godunov methods have been able to overcome this obstacle.

The interaction of electromagnetic fields with plasmas generates a wide range of complicated motions that do not appear in other states of matter. The large families of resultant waves fill the pages of entire books ([28], [29]). Even in the simplified approximation in which we operate, our plasmas contain two kinds

of sound waves and one kind of purely magnetic disturbance.

1.2.2 A Brief History

The computational plasma physics literature dates back to 1960. Most of the effort was concentrated in a few large institutions in the U.S., including Los Alamos and Lawrence Livermore National Laboratories, the Naval Research Laboratory in Washington and a handful of universities. The earliest MHD codes extended the capabilities of hydrodynamics codes that used finite difference methods to solve the fluid equations. One algorithm extended Harlow's hydrodynamic Particle in Cell variant [30] for use with MHD [31] and was used to model shock heating in plasmas.

The 1980s saw great progress in MHD simulations. Brio and Wu [32] developed an upwind finite difference method that allowed the treatment of magnetic discontinuities in shock tube problems. The Brio and Wu shock tube problem has since become a standard test for numerical MHD. Meanwhile, Evans and Hawley developed a Flux-Constrained Transport (FCT) method that exactly conserves the magnetic flux by expressing the magnetic induction with a discrete form of Stoke's Theorem [33]. This method produces magnetic fields that are manifestly divergence-free in a manner reminiscent of Yee's famous Finite Difference Time Domain (FDTD) method for treating electromagnetic waves [34]. This method is still being used and extended in modern codes.

In the last two decades, Godunov methods, which approximate solutions to Riemann's shock tube problem [27], were developed to accommodate conditions for magnetic fields ([35], [36]) as well as those for ordinary hydrodynamics. This is significant, since these methods can accommodate discontinuities more gracefully than methods based on spatial discretization of partial derivatives. Consequently, it has become common practice to compare a candidate method's test results to those of a high-order Godunov solver.

The development of the FCT method has continued to inspire finite difference and finite element schemes that produce divergence-free magnetic fields. Hyman and Shashkov describe a method for magnetic diffusion [37] based on discrete forms of equivalent theorems in vector calculus. Finite elements are also becoming more popular with the availability of more computing power, since they handle complicated

boundary conditions more gracefully. Rieben has applied the Vector Finite Element method, which also produces zero-divergence magnetic fields, to magnetic diffusion [38]. Variants of these finite difference and finite element methods are currently being developed to solve problems in resistive magnetohydrodynamics.

A number of modern MHD codes are now maintained by groups at universities and are available for use by interested scientists. ZEUS [39] is a family of 2D and 3D codes developed to study astrophysical flows. It implements Stone’s extensions to Evans and Hawley’s FCT algorithm and continues to be extended. Another astrophysical MHD code is Athena [40], which uses a high-order Godunov method in one, two, and three dimensions to study similar problems. We use Athena to verify some of our MHD test problems.

In the magnetic fusion community, NIMROD [41] is a major workhorse for studying plasmas in tokamaks and spheromaks. NIMROD uses two- and three-dimensional finite elements to solve the equations of resistive MHD in toroidal geometry. It has been used to do simulations of real experiments [42] and to examine several plasma configurations.

In spite of these efforts, simulating plasmoids remains difficult. The complicated and sometimes violent motions of the plasma causes problems for any method that represents space on a grid. On stationary grids, rapid motion can cause excessive numerical diffusion; on moving grids, nodes can cross one another, resulting in unphysical geometries.

1.3 Mesh-Free Methods

In this thesis we simulate plasmas on sets of computational points (*nodes*) that do not connect to each other to form a grid. Methods of this sort are most commonly called “mesh-free” methods and are characterized by the representation of spatial derivatives by “stencils” defined on sets of neighboring nodes. These methods enjoy a number of advantages over grid-based methods in problems with large deformations, since they do not rely on fixed node connectivity.

1.3.1 Smoothed Particle Hydrodynamics

SPH was developed independently by Gingold and Monaghan [43] for modeling aspherical stars and by Lucy [44] as a numerical technique for modeling binary star systems. Representing mass distributions with a disconnected set of points in space, the technique is well-suited to systems without obvious symmetries, or with symmetries that change significantly over the course of a simulation. Originally, SPH was understood in terms of the interpolation of a smooth function over a set of disordered nodes. The equations of fluid mechanics are very simply expressed in SPH as sums of nodal properties and derivatives of a smoothing kernel function that is evaluated on sets of interacting nodes. The nodes, which loosely represent fluid elements and which are assigned mass and move with a fluid velocity, are allowed to connect and reconnect with one another as necessary to maintain a continuum on which equations are solved.

The simplicity of the SPH formulation is striking—it is the same in one dimension as in three, for example. However, it often requires modification in order to avoid numerical instabilities. Some criticisms of the method are that it is inconsistent (unable to exactly reproduce functions of any class), is numerically unstable under tension, and unable to represent strong anisotropy. Fortunately, there exist variants of SPH that address each of these difficulties, as we will discuss in Chapter 3.

SPH is beginning to be used more by the engineering community for modeling fluid flows and solid mechanics. Its dynamic connectivity makes it suitable for problems that lack obvious symmetry. More recently, SPH-based diffusion methods have been used to study heat transfer ([4], [45]), radiation transfer ([46]), and magnetohydrodynamics ([47], [3]). In our research, we have extended SPH to model magnetic diffusion.

1.3.2 Mesh-Free Galerkin and Galerkin-Like Methods

Mesh-free methods reached the engineering community in the early 1990s when Nayroles [48] and Duarte [49] began developing finite-element-like methods that did not require fixed connectivity between nodes. Perhaps the best-known such approach is the Element-Free Galerkin (EFG) Method, formalized by

Belytschko [50]. These techniques use the same mathematical formalism as finite element methods, defining linear and bilinear forms in terms of shape functions amidst computational points called nodes. Instead of defining these shape functions in terms of polynomials confined to elements, the mesh-free methods use various approximations of functions on disjoint points, defining their own “mesh-free” shape functions. Several families of shape functions have since emerged and form the basis of corresponding mesh-free methods best characterized by the EFG method. All of these methods allow practitioners to set up problems in elasticity without using esoteric and costly software to generate complicated grids.

It is interesting to note that most of these methods rely on some notion of fixed connectivity in order to perform the numerical integration of the linear and bilinear forms that is characteristic of Galerkin methods. More precisely, these methods tend to use a fixed set of quadrature points on a background mesh in order to compute the integrals. Perhaps this is because the techniques were developed in pursuit of solid mechanics and not fluid mechanics. In this way, it may be claimed that the EFG method and its ilk are not truly “mesh-free.” Certainly, one can imagine the difficulty this may present for a computational domain that moves and deforms.

More recently, Atluri et. al have developed a family of Galerkin-like methods, called Meshless Local Petrov-Galerkin (MLPG) methods [5], that rely on no such background mesh. Rather, equations are solved on subdomains comprised of sets of interacting nodes and quadrature points that are defined dynamically within these subdomains. These kinds of methods continue to show great promise in solid mechanics, and are making inroads to fluid mechanics [51] and magnetic diffusion [52].

These finite-element-like mesh-free methods are important because they can treat *natural* boundary conditions, which are not easily expressed in SPH and other methods based on the “strong form” of their corresponding equations. In the context of magnetic fields, a natural boundary condition is a constraint on the curl of \mathbf{B} at the boundary, either in terms of the current density \mathbf{J} or an external electric field \mathbf{E} . It turns out that such boundary conditions are crucial to our application of interest, and so we explore a hybrid of the SPH and MLPG methods.

1.3.3 Least-Squares Approaches

Mesh-free methods usually fall into one of the above two categories, but there are other approaches, usually based on representations of functions by their Taylor series on discrete points involving the solution of a least-squares system. The Least Squares Finite Difference (LSFD) method [53] is an example of such a technique, and has been used for studies of heat transfer and hydrodynamics. The technique is simple, but suffers from two drawbacks. Functions represented by least-squares systems become less accurate when expressed on disordered points. In particular, when two points are close to one another, the equations for the Taylor series at these two points become “more linearly dependent” (an idea that separates discrete linear algebra from its continuous counterpart). Solutions to these linear systems are very inaccurate. Also, the second derivatives of functions expressed in this way tend to have significant oscillations, since it is not possible to guarantee the monotonicity of both a function and its derivatives. Nevertheless, this approach continues to be explored as a third approach to mesh-free methods.

In 2003 Maron and Howes [54] described a method for solving the equations of ideal magneto-hydrodynamics using a vector potential \mathbf{A} expressed as the sum of polynomial contributions whose coefficients are the solution of a least-squares system. This approach generates a magnetic field that is manifestly divergence-free, but the current density, which is the second derivative of \mathbf{A} , suffers from the same oscillations observed in the LSFD method. This approach seems not to have been explored recently.

More recently, Diyankov has proposed a flux-conserving mesh-free approach he calls the “uncertain grid method” for hydrodynamic flows [55]. This method solves the integral equations of hydrodynamics using a least-squares formulation, which allows fluxes to be well-defined. It is not clear how this method performs practically under disordered node distributions, but this “uncertain grid method” is a departure from the treatment of the differential hydrodynamic equations and deserves to be explored.

1.4 Features of the Method

The method we describe in this dissertation enjoys several advantages over more traditional methods.

- **Easy Problem Generation**

The most significant feature of mesh-free methods is that there is no mesh! Any computational scientist or engineer that has run a realistic simulation on a mesh knows that the mesh’s geometry can factor heavily into that simulation’s performance and accuracy, and also that generating complicated meshes is no easy task. Specialized software is often needed to determine nontrivial connectivity, and even the process of learning to use this software requires much time and determination. Mesh-free methods sidestep this entire process by deferring the process of connecting nodes until a simulation starts.

- **Conceptual Simplicity**

The SPH equations of motion use the *Lagrangian* formulation, in which the computational nodes follow the motion of the mass of the material therein. These equations are simpler than their *Eulerian* counterparts because the nonlinear advection term in the momentum equation is not present, and there is no need, a priori, to use flux limiting to maintain the monotonicity of signals. In addition, the treatment of multiple materials is greatly simplified, since each material can be treated as a separate set of nodes, and complicated interface tracking methods are unnecessary.

- **Robustness**

Mesh-based Lagrangian methods in fluid mechanics run into trouble when the elements within the mesh become sufficiently deformed. When nodes cross, element volumes can become negative, rendering meaningless solutions to the equations of motion. Even before the mesh “tangles” in this way, extreme aspect ratios of elements can decrease the accuracy of a simulation. Such difficulties are often addressed by a periodic remapping of the nodes to a more “regular” configuration, at which point the benefit of the Lagrangian

description (the elimination of the nonlinear advection term) is diminished. SPH does not produce such troublesome geometries, as its connectivity is completely dynamic. While the accuracy of a simulation may suffer under large amounts of node disorder, no periodic remapping is required to keep a simulation running.

- **Galilean Invariance**

All of the equations we represent are expressed in vector form, and thus do not implicitly or explicitly depend on any given coordinate system. Thus, as we change the distribution of computational points, we do not expect the results of simulations to change significantly. We have found this to be largely the case, though some care must be used in placing boundary nodes in the vicinity of non-trivial boundaries.

- **Adaptive Resolution**

A distinguishing feature of SPH is its coupling of resolution to mass. The volume of an SPH fluid element is expressed as $\Delta V = m/\rho$, where m is the mass of that element and ρ is its mass density. This means that the regions with the most mass also have the highest resolution. For many problems, this is a welcome relationship. Further, it is relatively easy to insert new nodes into a simulation if more resolution is needed, either by inflow at a boundary (which we discuss in Chapter 3) or by splitting nodes (which we do not discuss but which has been explored by a few authors ([56], [57])). The use of other mesh-free methods in conjunction with SPH does not undermine either of these features.

- **Easier Parallel Computing**

Unlike many grid methods, SPH can be parallelized relatively easily using the same method that it uses for enforcing boundary conditions. In addition, domain decomposition can easily be done initially and for load balancing during simulations. Parallel mesh generation, on the other hand, can be complicated or limited. In structured grids, domain decomposition must often be done along the lines of the mesh. In block-structured grids, parallelism is usually done in terms of blocks, which makes decomposition awkward for certain numbers of processors. In unstructured grids, domain decomposition can be very difficult and requires a lot of bookkeeping.

1.4.1 Limitations

Obviously, no numeric method is perfect, and SPH brings with it a number of challenges shared by many other mesh-free methods.

- **Lack of Topological Constraints**

Since SPH quantities are all node-centered, it is difficult to imagine how to formulate so-called compatible spatial differential operators. This is perhaps the most significant drawback of SPH in treating magnetic fields, since we cannot guarantee exact magnetic flux conservation or divergence-free \mathbf{B} fields, in contrast to grid-based techniques like that of Flux-Constrained Transport.

Some authors have taken some rudimentary steps toward more sophisticated mesh-free topologies. For example, Dual Particle Dynamics [58] is a mesh-free method that supports staggered quantities by introducing “stress” points in addition to the usual “material” points, but the more nebulous definition of resolution in mesh-free methods makes it difficult to express differential operators in a manifestly compatible way.

- **Computational Expense**

SPH is quite expensive compared to any grid-based method with fixed connectivity, since it dynamically determines its connectivity one or more times during each time step. In principle, this reconnection process is similar to those used in *free-Lagrange* mesh algorithms, in which an unstructured mesh that becomes tangled reconnects its nodes to achieve a smoother topology. Fortunately, the simple node-to-node simplicity of mesh-free connectivity makes the reconnect cheaper in SPH than in these mesh-based schemes.

- **Difficult Boundary Treatment**

Boundary conditions require somewhat special attention in SPH compared with Galerkin-like methods, since the boundary terms are usually neglected in its formulation. We will discuss the ways in which we address boundary conditions near straight and curved surfaces and at free boundaries.

1.5 Dissertation Roadmap

In Chapter 2 we discuss the physics necessary for modeling the plasmas we have mentioned. We introduce the equations of resistive magnetohydrodynamics (MHD) as an approximation to the dynamics of plasmas in magnetic fields, exploring their explanatory power and their limitations in this respect. We examine the behavior of magnetic fields in different materials and in vacuum, and we explore different models of electrical conductivity that govern this behavior. In Chapter 3 we present the Smoothed Particle Hydrodynamics approximation and discuss its various forms. We add effects from magnetic fields in Chapter 4, extending existing methods to model magnetic diffusion in resistive materials like laboratory plasmas. Both of these chapters include test problems that verify the various aspects of the method, and we perform rigorous convergence studies to show that the method obtains convergent solutions to these problems. Then, in Chapter 5, the method is applied to a plasma compression experiment performed by the CTIX group, and we show that there is good qualitative agreement between the gathered diagnostics and the computational predictions. We conclude the thesis in Chapter 6, summarizing the method's performance and discussing future directions for improvement.

Chapter 2

The Physics of Magnetized Plasmas

In this chapter, we give a brief overview of how plasmas are analyzed mathematically. We motivate and develop the magnetohydrodynamic approximation of a magnetized plasma and justify its use for the problems that we will discuss in later chapters. We use MKS units and the corresponding equations of electrodynamics.

2.1 Charged Particles in Magnetic Fields

In the most general sense, a plasma is a collection of unbound charged particles. A plasma is created when a gas is heated to the point at which the electrons separate from the ions but remain in the vicinity so that positive and negative charges are allowed to interact with one another. There are several methods used to create plasmas in the laboratory. One method is simply to hit a material with a high-powered laser; another is to run a large current through a gas.

In a fully ionized gas, there is tremendous electrostatic repulsion between the ions. Plasmas are fundamentally different from this because the constituent electrons and ions act collectively to neutralize excess electric charge. A positively-charged ion, as it moves through the plasma, tends to gather a negatively-charged cloud of electrons around itself. This cloud shields the charge of the ion so that other ions outside

of a radius λ_D , called the *Debye length*, are not repelled by it. This effect is called *Debye shielding*, and is crucial for the very existence of plasmas. In fact, we can define a plasma as an ionized gas that has many particles in a cube with sides of length λ_D .

Plasmas are rich in their dynamics and their variety, and we make no attempt at any comprehensive survey in this dissertation. We concern ourselves with plasmas in strong magnetic fields that have roughly equal numbers of positive and negative charges (no net charge) and describe some basic effects that are necessary to understand the essential physics. We begin with a brief look at the dynamics of charged particles in magnetic fields.

2.1.1 Plasma oscillations

Consider a particle with mass m and electric charge q moving with velocity \mathbf{v} in the presence of electric and magnetic fields \mathbf{E} and \mathbf{B} . The force on the particle is

$$\mathbf{F} = m \frac{d\mathbf{v}}{dt} = q[\mathbf{E} + \mathbf{v} \times \mathbf{B}] \quad (2.1)$$

The electric field pushes the particle in a straight line while the magnetic field produces a force that is perpendicular to \mathbf{v} . If \mathbf{B} is uniform within some region, we can consider this a centripetal force that moves the particle in a circle

$$\mathbf{F}_c = \frac{mv^2}{r} \hat{r} = q\mathbf{v} \times \mathbf{B} \quad (2.2)$$

where r is the radius of the circle traced by the particle in the plane whose normal vector is parallel to \mathbf{v} . We can express the motion of the particle in terms of its angular frequency $\Omega = v/r$, which, according to (2.2), is

$$\omega_c = \frac{|q||\mathbf{B}|}{m} \quad (2.3)$$

The circular motion of electrons in the presence of a magnetic field is called *cyclotron motion*, and Ω is the *cyclotron frequency*. The result is that, in regions of relatively constant \mathbf{E} and \mathbf{B} , the electrons follow a helical

motion in the direction of \mathbf{E} with orbits that are confined to magnetic field lines within a radius depending on the strength of the magnetic field. This picture of electron currents in magnetic fields is important to keep in mind when we attempt to simplify the equations of motion later in this chapter. The electrons are also attracted to the positively-charged ions in the plasma, and oscillate about them with a characteristic *electron plasma frequency*

$$\omega_p = \sqrt{\frac{n_0 e^2}{\epsilon_0 m_e}}, \quad (2.4)$$

where n_0 is the average number density of the electrons, m_e is the electron mass, and e is the electronic charge.

Everything we have discussed for the electrons in a plasma holds for the ions as well—there is indeed an ion cyclotron frequency and an ion plasma frequency. Since these frequencies depend inversely upon the mass of their particles, the ion frequencies are much smaller than those of electrons, and so these motions occur at far slower time scales. This observation is hardly surprising if one considers an equivalent force $F = ma$ applied to particles of very different masses, but it is worth mentioning because the fact that the electrons move at much faster time scales than the ions is the basis for the magnetohydrodynamic (MHD) approximation. In MHD, these motions are represented by algebraic relationships between the electrons, the ions, and the magnetic field. For example, magnetic field lines capture electrons, producing currents proportional to the curl of \mathbf{B} , as we shall see later. Similarly, the electron-ion plasma interactions tend to neutralize the electric field over larger timescales, resulting in a generalized form of Ohm's Law.

We will now attempt to provide a more precise picture of the motion of particles in a plasma.

2.2 Mathematical description

The most straightforward way to describe a classical system of interacting charged particles is to use the Lorenz force equation (2.1) in concert with Maxwell's equations. That is, the force on a particle i at position \mathbf{x}_i with velocity \mathbf{v}_i in a system of N particles is

$$\mathbf{F}_i = q_i[\mathbf{E}(\mathbf{x}_i) + \mathbf{v}_i \times \mathbf{B}(\mathbf{x}_i)] \quad (2.5)$$

where \mathbf{E} and \mathbf{B} are the electric and magnetic fields in the region. These fields evolve in time according to Maxwell's Equations:

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \quad (2.6)$$

$$\nabla \cdot \mathbf{E} = \frac{\rho_e}{\epsilon} \quad (2.7)$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \frac{\partial(\epsilon \mathbf{E})}{\partial t} \quad (2.8)$$

$$\nabla \cdot \mathbf{B} = 0 \quad (2.9)$$

where ρ_e is the electric charge density, \mathbf{J} is the electric current density, ϵ is the permittivity of the plasma, and μ_0 is the permeability of free space in vacuum.

If we express ρ_e and \mathbf{B} in terms of Dirac delta functions, the fields can be written in terms of Green's functions and we can solve these continuum equations in terms of the discrete system of particles, whereupon we may apply the forces to the particles and repeat the procedure. However, this method of solution is clearly not scalable. The plasmas we are interested typically contain around 10^{15} particles per cubic meter, so N is monstrous!

Plasmas as Non-magnetic Materials

Before we continue, we should note that we have chosen to avoid mentioning the “auxiliary” fields \mathbf{H} and \mathbf{D} in our treatment because it is common practice to discuss plasma physics in terms of the fundamental fields \mathbf{B} and \mathbf{E} . One of the reasons for this is that plasmas do not behave like linear magnetic materials. Chen shows very clearly that if one computes the net magnetization vector \mathbf{M} of a plasma, this quantity varies as the *inverse* of \mathbf{B} , precluding any linear relationship between \mathbf{B} and \mathbf{H} [21].

It is true that plasmas can behave like dielectrics and that the use of \mathbf{D} can be justified, but this turns out to be irrelevant to the approximation that we will use to describe our magnetized plasmas.

2.2.1 Plasma Kinetic Theory

Those tasked with describing systems with monstrous values of N developed a technique for replacing the discrete force law (2.1) with a differential equation defined in a continuous space in the hope that mathematical methods would be of some service. In this picture, a plasma consists of a set of *species* $\{s_i\}$ made up of particles of mass m_s and charge q_s . For each species, we define a *distribution function* $f_s(\mathbf{x}, \mathbf{v}, t)$ such that the number of particles in species s located within the region $\Omega_{\mathbf{x}} = [x_1, x_2] \otimes [y_1, y_2] \otimes [z_1, z_2]$ and having a velocity within the range $\Omega_{\mathbf{v}} = [v_{x1}, v_{x2}] \otimes [v_{y1}, v_{y2}] \otimes [v_{z1}, v_{z2}]$ at time t is

$$N_s(t) = \int_{\Omega_{\mathbf{x}}} \int_{\Omega_{\mathbf{v}}} f(\mathbf{x}, \mathbf{v}, t) d\mathbf{x} d\mathbf{v} \quad (2.10)$$

The total number of particles in the plasma is just the sum of the number of particles of each species, which should remain constant in the absence of particle creation/annihilation:

$$N = \sum_s N_s(t) \quad (2.11)$$

Thus, our distribution function f is a description of the plasma in *phase space*.

To solve the Maxwell's equations (2.7) and (2.8) we need expressions for ρ_e and \mathbf{J} . These are averages of the charge distributions and their currents across all species and are given by

$$\rho_e(\mathbf{x}, t) = \sum_s q_s \int f(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \equiv \sum_s q_s n_s(\mathbf{x}, t) \quad (2.12)$$

$$\mathbf{J}(\mathbf{x}, t) = \sum_s q_s \int \mathbf{v} f(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \quad (2.13)$$

where we have integrated over all of velocity space and introduced the *number density* n_s , which is the number of particles of species s per unit volume:

$$n_s(\mathbf{x}, t) \equiv \int f(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \quad (2.14)$$

The distribution function f evolves in time according to a continuity equation wrought by our assumption that particles are neither created nor destroyed:

$$\frac{df}{dt} \equiv \frac{\partial f_s}{\partial t} + \mathbf{v} \cdot \frac{\partial f_s}{\partial \mathbf{x}} + \mathbf{a}_s \cdot \frac{\partial f_s}{\partial \mathbf{v}} = \Pi_s \quad (2.15)$$

The term on the right hand side is the rate of change in f_s due to collisions and interactions with other species. The other terms are “convective” terms in phase space. \mathbf{a}_s is the force per unit mass exerted by electromagnetic interactions according to the Lorentz force law (2.1). A more familiar form of (2.15) is

$$\frac{df}{dt} \equiv \frac{\partial f_s}{\partial t} + \mathbf{v} \cdot \nabla f_s + \frac{q_s}{m_s} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f_s}{\partial \mathbf{v}} = \Pi_s \quad (2.16)$$

This harmless-looking continuity equation is called the Boltzmann equation, and when combined with Maxwell’s equations ((2.6), (2.7), (2.8), (2.9)) exactly describes a plasma in terms of a continuum instead of an enormous system of particles. However, it is folly to think that this system of equations is any more practical to solve exactly than our original set! In its exact form, each distribution function f is still a set of Dirac delta functions in a six-dimensional space, and the fields are still a superposition of Green’s functions that produce very noisy fields.

Nevertheless, this is a better point of departure. There are several approximations of the Boltzmann equation, each emphasizing a different area of interest in plasma physics. When $\Pi_s = 0$, the plasma is said to be collisionless and (2.16) is called the *Vlasov Equation* and is used in laser plasma physics, where collisions are relatively unimportant. There are various models that treat collisionality by adding terms to this equation. Nicholson [59] and [21] discuss the kinetic theory of plasmas in more detail. We are concerned with strongly magnetized plasmas that are dense enough to be dominated by collisions, so we can make simplifying assumptions about the distribution functions $\{f_s\}$ that allow us to reduce the kinetic description to a much more practical fluid description.

2.2.2 Fluid Equations

In plasmas where collisions are frequent enough for the particles to begin to equilibrate with one another, we assume that system is in *local thermal equilibrium* (LTE). That is, each particle is in approximate

thermal equilibrium with its neighbors, though not necessarily with particles outside direct contact. In LTE, we assume that each distribution function is Maxwellian, i.e. that

$$f_s(\mathbf{x}, \mathbf{v}, t) = g_s(\mathbf{x}, t) \exp(-|\mathbf{v} - \mathbf{v}_{0s}|^2/a^2) \quad (2.17)$$

This configuration represents a Gaussian distribution of velocities about an average “bulk” velocity of the species s and allows us to halve the number of dimensions in the problem. Instead of working with distribution functions in phase space (\mathbf{x}, \mathbf{v}) , we define the problem in terms of fields associating quantities (bulk velocity \mathbf{v}_s , mass density ρ_s , magnetic field \mathbf{B} , and so forth) with points in three-dimensional space. Note that $g_s(\mathbf{x}, t)$ must be normalized appropriately for the distribution function.

This simplification is described in careful detail by Braginskii [60] and several others ([59], [21]). The distribution function is multiplied by the zeroth, first, and second powers of the velocity and integrated over all velocity space to produce its first three *velocity moments*

$$I_{s0}(\mathbf{x}, t) \equiv \int f_s(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \quad (2.18)$$

$$\mathbf{I}_{s1}(\mathbf{x}, t) \equiv \int \mathbf{v} f_s(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \quad (2.19)$$

$$\mathbf{I}_{s2}(\mathbf{x}, t) \equiv \int \mathbf{v} \mathbf{v} f_s(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \quad (2.20)$$

Note that the zeroth velocity moment of the species, I_{s0} , is just its number density n_s (2.14). The vector-valued first moment \mathbf{I}_{s1} is the momentum per unit mass $n_s \mathbf{v}_s$, and the trace of the second moment tensor \mathbf{I}_{s2} is the specific energy (energy per unit mass) u_s .

Integrating the first three moments of the Boltzmann function equation for the s th species (2.15) over all velocity space produces conservation equations for I_{s0} , \mathbf{I}_{s1} , and $\text{Tr} \mathbf{I}_{s2}$ that can be rearranged in terms of the quantities n_s , \mathbf{v}_s , and u_s . The integration of the first and second moments requires an expression for the plasma’s pressure, and the assumption of LTE allows us to use the ideal gas law $p_s = n_s k_B T$, where T is the temperature of the plasma. For a simple 2-species plasma involving one kind of ion (written with subscript i) and the electrons (subscript e), the integrated moment equations produce the fluid equations

$$\frac{\partial n_i}{\partial t} + \nabla \cdot (n_i \mathbf{v}_i) = 0 \quad (2.21)$$

$$m_i n_i \left[\frac{\partial \mathbf{v}_i}{\partial t} + (\mathbf{v}_i \cdot \nabla) \mathbf{v}_i \right] = e m_i n_i (\mathbf{E} + \mathbf{v}_i \times \mathbf{B}) - \nabla p_i + \mathbf{P}_i \quad (2.22)$$

$$\frac{\partial u_i}{\partial t} + \nabla \cdot (u_i \mathbf{v}_i) = 0 \quad (2.23)$$

$$\frac{\partial n_e}{\partial t} + \nabla \cdot (n_e \mathbf{v}_e) = 0 \quad (2.24)$$

$$m_e n_e \left[\frac{\partial \mathbf{v}_e}{\partial t} + (\mathbf{v}_e \cdot \nabla) \mathbf{v}_e \right] = -e m_e n_e (\mathbf{E} + \mathbf{v}_e \times \mathbf{B}) - \nabla p_e + \mathbf{P}_e \quad (2.25)$$

$$\frac{\partial u_e}{\partial t} + \nabla \cdot (u_e \mathbf{v}_e) = 0 \quad (2.26)$$

These equations, called the ion and electron plasma fluid equations, resemble the ordinary equations of fluid mechanics without viscosity. This is not the only set of fluid equations one can derive for a plasma, but it is the one consistent with the assumption of local thermal equilibrium. Hazeltine [61] discusses the more general moment equations and the definition of the plasma pressure.

It is somewhat amazing that electrons and ions in plasmas can be treated like interpenetrating fluids, given the strength of the electromagnetic force. But one must keep in mind that Debye shielding keeps the electromagnetic force small. Note that the electron plasma frequency (2.4) can be reproduced by these equations if one assumes infinitesimal wave solutions $e^{i\omega_p t}$ for the density and the velocity in a uniform plasma.

2.3 The Magnetohydrodynamic Approximation

We have assumed that the particles in a collision-dominated plasma interact with one another in such a way to produce regions of local thermal equilibrium. This equilibrium occurs on a time scale that is much greater than that of electromagnetic wave propagation. In fact, since the ions carry most of the kinetic energy in the plasma because their greater mass, the motion of electrons also occurs at a smaller time scale than this equilibration process. Both of these observations lead to simplifications in Maxwell's Equations.

2.3.1 The Induction Equation

Since we are not interested in electromagnetic waves, we neglect the displacement current term in Ampere's Law to obtain a simple relationship between \mathbf{J} and \mathbf{B} :

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} \quad (2.27)$$

Further, because electrons move at faster time scales than we can capture, we cannot represent the corresponding electrostatic oscillations in the plasma, and the electric field cannot be determined using Gauss's Law (2.7). In fact, the electric field inside a perfectly conducting plasma will be zero in the reference frame of the plasma, since the electrons can move “instantaneously” to neutralize the electrostatic force:

$$\mathbf{E}' = \mathbf{0} \quad (2.28)$$

In an imperfect conductor, \mathbf{E} can be related to \mathbf{J} using the electrical conductivity of the plasma:

$$\mathbf{E}' = \sigma^{-1} \mathbf{J} \quad (2.29)$$

This is Ohm's Law. The inverse of the electrical conductivity σ is often written as the electrical resistivity $\eta \equiv \sigma^{-1}$. If we transform \mathbf{E} to a stationary reference frame (following e.g. [62]), we obtain

$$\mathbf{E} = -\mathbf{v} \times \mathbf{B} + \eta \mathbf{J} \quad (2.30)$$

where we have introduced an unsubscripted *bulk velocity* \mathbf{v} representing the density-weighted average of the velocities of the various species in the plasma:

$$\mathbf{v} \equiv \frac{\sum_{s=1}^N n_s \mathbf{v}_s}{\sum_{s=1}^N n_s} \quad (2.31)$$

When \mathbf{E} is thus prescribed, Faraday's Law reduces to

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} = \nabla \times (\mathbf{v} \times \mathbf{B}) - \nabla \times (\eta \mathbf{J}) - \nabla \times \mathbf{E}_{ext} \quad (2.32)$$

where \mathbf{E}_{ext} is any specified external electric field. Eq. (2.27) lets us rewrite the \mathbf{J} term as

$$\nabla \times (\eta \mu_0^{-1} \nabla \times \mathbf{B})$$

and if we assume for the moment that η is spatially homogeneous, we can use the vector identity

$$\nabla \times \nabla \times \mathbf{A} = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2 \mathbf{A} \quad (2.33)$$

and (2.9) to rewrite (2.32) as

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{v} \times \mathbf{B}) - \frac{\eta}{\mu_0} \nabla^2 \mathbf{B} - \nabla \times \mathbf{E}_{\text{ext}} \quad (2.34)$$

We have reduced the field equations to the Ohm's Law (2.30) and the equation for \mathbf{B} , (2.34), which is known as the *induction equation*. When we use the induction equation to evolve \mathbf{B} , we can obtain \mathbf{J} using (2.27). Evidently there is no need to keep track of charge densities or electrical currents directly.

In reality, the form of η will depend on the state of the plasma itself. We will refrain from discussing specific resistivity models at this point and revisit this topic when we discuss magnetic reconnection later in this chapter.

A special case of (2.34) worth mentioning is the *magnetic diffusion equation*, in which $\mathbf{v} = \mathbf{0}$:

$$\frac{\partial \mathbf{B}}{\partial t} = \frac{\eta}{\mu_0} \nabla^2 \mathbf{B} \quad (2.35)$$

This equation gets its name from its similarity (superficial and actual) to the heat equation. The essential idea is that a magnetic field in an imperfect conductor will diffuse out of the conductor at some rate proportional to its electrical resistivity. By analogy to a thermal conductivity, the ratio η/μ_0 is sometimes called the *magnetic diffusivity*. This equation is important for modeling stationary conducting materials in the vicinity of a magnetized plasma. We will discuss this equation and its progenitor (2.34) in more detail later in this chapter; its numerical treatment appears in Chapter 4.

2.3.2 The Single-Fluid Equations

It remains for us to simplify the fluid equations that we obtained from the velocity moments of the distribution function. To reduce the electron and ion fluid equations to a single fluid, we define bulk quantities for the mass density ρ , velocity \mathbf{v} , and current density \mathbf{J} , using the quasi-neutral approximation $n_i \approx n_e = n$:

$$\rho \equiv n_i m_i + n_e m_e \approx n(m_i + m_e) \quad (2.36)$$

$$\mathbf{v} \equiv \frac{1}{\rho} (n_i m_i \mathbf{v}_i + n_e m_e \mathbf{v}_e) \approx \frac{m_i \mathbf{v}_i + m_e \mathbf{v}_e}{m_i + m_e} \quad (2.37)$$

$$\mathbf{J} \equiv e(n_i \mathbf{v}_i - n_e \mathbf{v}_e) \approx ne(\mathbf{v}_i - \mathbf{v}_e) \quad (2.38)$$

After rearranging (2.21) - (2.26) in terms of these bulk quantities, we end up with the following single-fluid conservation equations:

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \mathbf{v}) \quad (2.39)$$

$$\frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} = -\frac{1}{\rho} \nabla p + \frac{\mathbf{J} \times \mathbf{B}}{\rho} \quad (2.40)$$

$$\frac{\partial u}{\partial t} = -\nabla \cdot (u \mathbf{v}) \quad (2.41)$$

Here, we have also used a bulk specific thermal energy u that tallies the total energy per unit mass of the single fluid. These and (2.34) are the *resistive MHD equations* and form the basis for the plasma simulations in this dissertation. The *ideal MHD equations*, of interest in astrophysics, are obtained when one sets the resistivity η to 0 in the induction equation.

Other derivations of the MHD equations are possible. Biskamp [63] and Jackson [62] both derive these equations from simpler arguments. Davidson [9] starts with the equations of fluid mechanics and considers the extra terms necessary when a conducting fluid is placed in a magnetic field. This approach suggests that MHD works for modeling a plasma to the extent that the plasma behaves like a conducting fluid.

2.3.3 Boundary Conditions at Material Interfaces

In a laboratory setting, a plasma inevitably interacts with its surrounding environment: plasmas are notorious for their instabilities even when placed in large magnetic fields. Biskamp [63], Nicholson [59], and Chen [21] discuss various *kink* and *sausage* instabilities arising from the $\mathbf{J} \times \mathbf{B}$ force, which remains the largest obstacle to controlled fusion in the laboratory.

The presence of boundaries, particularly conducting walls, is a significant factor in the behavior of plasmas. If a plasma comes into contact with a conducting wall, we may represent the other conductor as a stationary wall in which the field's behavior is dictated by the magnetic diffusion equation (2.35). Here, η becomes a function of space, since the plasma has a certain electrical resistivity and the wall has another. If the wall is a very good conductor, the transverse magnetic field penetrates only into a very shallow region [62] and any induced currents run along the surface. In this case, the relationship between the *surface current* \mathbf{K} (amount of charge per unit transverse length along the surface) and the field \mathbf{B} is

$$\mathbf{n} \times \mathbf{B} = \mathbf{K} \quad (2.42)$$

where \mathbf{n} is the normal vector of the surface. With the vector identity $\mathbf{n} \times (\mathbf{n} \times \mathbf{B}) = \mathbf{n}(\mathbf{n} \cdot \mathbf{B}) - \mathbf{B}(\mathbf{n} \cdot \mathbf{n})$, the equation for \mathbf{B} at the interface becomes

$$\mathbf{B} = B_n \mathbf{n} - \mathbf{n} \times \mathbf{K} \quad (2.43)$$

Notice that the normal component B_n of \mathbf{B} is continuous across the interface. This makes sense, since any normal discontinuity would produce a configuration for which $\nabla \cdot \mathbf{B} \neq 0$, in violation of (2.9).

We must also examine the interaction of a plasma with vacuum in the case that the vacuum contains an electromagnetic field, since external electric fields can drive currents that generate magnetic fields. In the vacuum region, the magnetic field is given by a magnetostatic configuration in which $\nabla \times \mathbf{B} = \mathbf{0}$, since there are no charge carriers in the vacuum to support currents. However, the surface of interaction between the plasma and the vacuum can contain current which means that $\nabla \times \mathbf{B} \neq \mathbf{0}$ there. If we define a vector potential

\mathbf{A} such that $\mathbf{B} = \nabla \times \mathbf{A}$, then the magnetic field in the vacuum may be obtained by solving the boundary value problem

$$\nabla \times \nabla \times \mathbf{A} = \mathbf{0} \quad (2.44)$$

with $\nabla \times \mathbf{A} = \mathbf{B}$ on the boundary, where \mathbf{B} is the value of the field in the plasma. Unfortunately, since \mathbf{B} is evolved using (2.34) in the plasma, it is often difficult in practice to couple the magnetohydrodynamic and magnetostatic solutions together at this interface.

A solution used by many in the field ([41], for example) is to pretend that the vacuum is simply a very poor conductor, to assign it a relatively large value for η , and to treat it like a conducting wall as described above. This approach presents various computational problems which we will discuss in due time, but it is a simple method that allows the treatment of these important interfaces.

2.3.4 Lagrangian Formulation

Up to this point, we have discussed the plasma in terms of quantities that are measured from a stationary reference frame. These equations are expressed in their *Eulerian form* in the language of hydrodynamics. Suppose now that we wish to use a coordinate system that moves with the material. If we define a *convective derivative* $\frac{D}{Dt} \equiv \partial/\partial t + \mathbf{v} \cdot \nabla$ that absorbs the convective term in the momentum equation, the resistive MHD equations ((2.39), (2.40), (2.41), (2.34)) may be written in their *Lagrangian form* [61]:

$$\frac{D\rho}{Dt} = \rho \nabla \cdot \mathbf{v} \quad (2.45)$$

$$\frac{D\mathbf{v}}{Dt} = -\frac{\nabla p}{\rho} + \frac{\mathbf{J} \times \mathbf{B}}{\rho} \quad (2.46)$$

$$\frac{Du}{Dt} = \frac{p}{\rho} \nabla \cdot \mathbf{v} + \frac{\eta J^2}{\rho} \quad (2.47)$$

$$\frac{D\mathbf{B}}{Dt} = \mathbf{B} \nabla \cdot \mathbf{v} - \mathbf{B} \cdot \nabla \mathbf{v} + \frac{\eta}{\mu_0} \nabla^2 \mathbf{B} - \nabla \times \mathbf{E}_{ext} \quad (2.48)$$

Since we simulate plasma with moving sets of points representing fluid elements, the Lagrangian formula-

tion of the MHD equations is more appropriate for our purposes than the Eulerian one. For ordinary fluid dynamics problems, the Lagrangian viewpoint is often simpler, as it eliminates a nasty convective term in the momentum equation. For MHD, this advantage is paired with a drawback in the more complicated-looking induction equation (2.48). Nevertheless, this form of the equations allows us to define ρ , \mathbf{v} , u , and \mathbf{B} on sets of moving points without using dissipative remapping processes.

2.4 The Physics of MHD

2.4.1 MHD Waves

In ordinary hydrodynamics, compressional sound waves propagate through a fluid at a characteristic *sound speed* c_s . This speed is defined as the coefficient that corresponds to the speed c of a solution to the wave equation $\partial u / \partial t = c \partial u / \partial x$ representing the linearized momentum equation [64]:

$$c_s \equiv \sqrt{\frac{\partial p}{\partial \rho}} \quad (2.49)$$

By contrast, we can expect to see *three* different kinds of waves in a plasma described by MHD: two compressional waves similar to the ordinary sound waves we have mentioned above, and one longitudinal wave that has no analog in ordinary fluid mechanics. We can best understand these waves by re-examining the forces acting on the plasma. The right-hand side of the momentum equation is

$$-\frac{\nabla p}{\rho} + \frac{\mathbf{J} \times \mathbf{B}}{\rho} \quad (2.50)$$

The first term is responsible for the purely hydrodynamic sound waves. We can rewrite the second term using (2.27) and a vector identity:

$$\frac{\mathbf{J} \times \mathbf{B}}{\rho} = \frac{1}{\rho \mu_0} (\nabla \times \mathbf{B} \times \mathbf{B}) \quad (2.51)$$

$$= \frac{1}{\rho \mu_0} \left[\frac{\nabla(B^2)}{2} + \mathbf{B} \cdot \nabla \mathbf{B} \right] \quad (2.52)$$

In homage to solid mechanics, the quantity

$$\mathbf{T} = \frac{1}{\mu_0} \left(\mathbf{B}\mathbf{B} - \frac{1}{2} B^2 \mathbf{1} \right) \quad (2.53)$$

(where $\mathbf{1}$ is the 3×3 identity tensor) is called the *Maxwell stress tensor*. In this language the $\mathbf{J} \times \mathbf{B}$ force is written $\nabla \cdot \mathbf{T}$, or

$$\frac{1}{\mu_0} \left(\frac{\nabla B^2}{2} + \mathbf{B} \cdot \nabla \mathbf{B} \right) \quad (2.54)$$

The first term is identical in form to the pressure force in ordinary hydrodynamic flows, with a “magnetic pressure” equal to $B^2/2\mu_0$. This means that magnetic field lines tend to repel one another. Using this concept of magnetic pressure, we can determine a “magnetosonic” speed for compressional MHD waves. Here there are two possible speeds instead of one, since ∇p and $\nabla B^2/2\mu_0$ may either be aligned or opposed to each other. Regions in which these gradients are aligned contain *fast magnetosonic waves*, and those in which gradients are opposed contain *slow magnetosonic waves*. The characteristic speeds of these waves are

$$v_{\pm} = \left\{ \frac{1}{2} \left[v_A^2 + c_s^2 \pm \sqrt{(v_A^2 + c_s^2)^2 - 4v_A^2 c_s^2 \cos^2 \theta} \right] \right\}^{1/2} \quad (2.55)$$

where v_+ is the fast magnetosonic speed, v_- is the slow magnetosonic speed, and v_A is the *Alfvén speed* defined by

$$v_A = \frac{|\mathbf{B}|}{\sqrt{\mu_0 \rho}} \quad (2.56)$$

Magnetosonic waves behave much like ordinary sound waves in that disturbances in the mass density, fluid velocity, and magnetic field are all propagated at the same speed.

The second term in (2.54) is a “magnetic tension” force. This force tends to restore curved \mathbf{B} field lines to a straight configuration, much like the tension in a plucked guitar string returns it to its equilibrium position. Any disturbance in the curvature of a magnetic field will propagate at the Alfvén speed, and disturbances of this sort are called *Alfvén waves* [65]. Unlike the aforementioned magnetosonic waves, Alfvén waves are purely magnetic disturbances.

Chen [21], Jackson [62], and Biskamp [63] all discuss magnetosonic waves and Alfvén waves in greater detail than we have here.

Shocks

Because the MHD equations are nonlinear like the equations of fluid mechanics, small magneto-acoustic disturbances in a fluid may grow larger over time, steepening into shock waves similar to those found in non-magnetized fluids ([66]). The fast and slow magnetosonic waves we have discussed can thus develop into fast and slow shocks. The conservation equations near these shocks are much more rich than the relations for ordinary sound waves since there are multiple wave speeds involved; considerable effort has been put forth to study this problem ([36], [67], [68], [69]). Figure 2.1 shows a train of waves produced by a magnetized shock tube described by Brio and Wu [32]. This problem illustrates well the various waves we have discussed, and has become a standard MHD test problem. We analyze it further in Chapter 4.

2.4.2 The Behavior of the Magnetic Field

The magnetic field evolves according to the induction equation, which for the moment we will analyze in its Eulerian form (2.34). The first term in the equation describes the movement of the magnetic field with the motion of the material, and the second term describes the diffusion of the field through the medium. The third term is included for open systems that include external fields which induce electrical currents at the boundary.

We derived this equation from Lenz’s Law ((2.6)), which states that a change in magnetic flux $\Phi_B \equiv \int \mathbf{B} \cdot \mathbf{n} dS$ produces an electromotive force that acts to counter that change [62]. The first term in (2.34) expresses the conservation of magnetic flux within a conducting material. This is evident if one neglects magnetic diffusion and integrates $\mathbf{n} \cdot (\mathbf{v} \times \mathbf{B})$ over the surface of the material to compute $\partial\Phi/\partial t$: the integral is zero [63]. In perfect conductors, this conservation is exact. Thus, if a perfectly conducting material compresses, its magnetic field intensifies to conserve the magnetic flux; if the material expands, the field relaxes to a lower intensity. In fact, one can follow the evolution of the field by etching the field lines into the

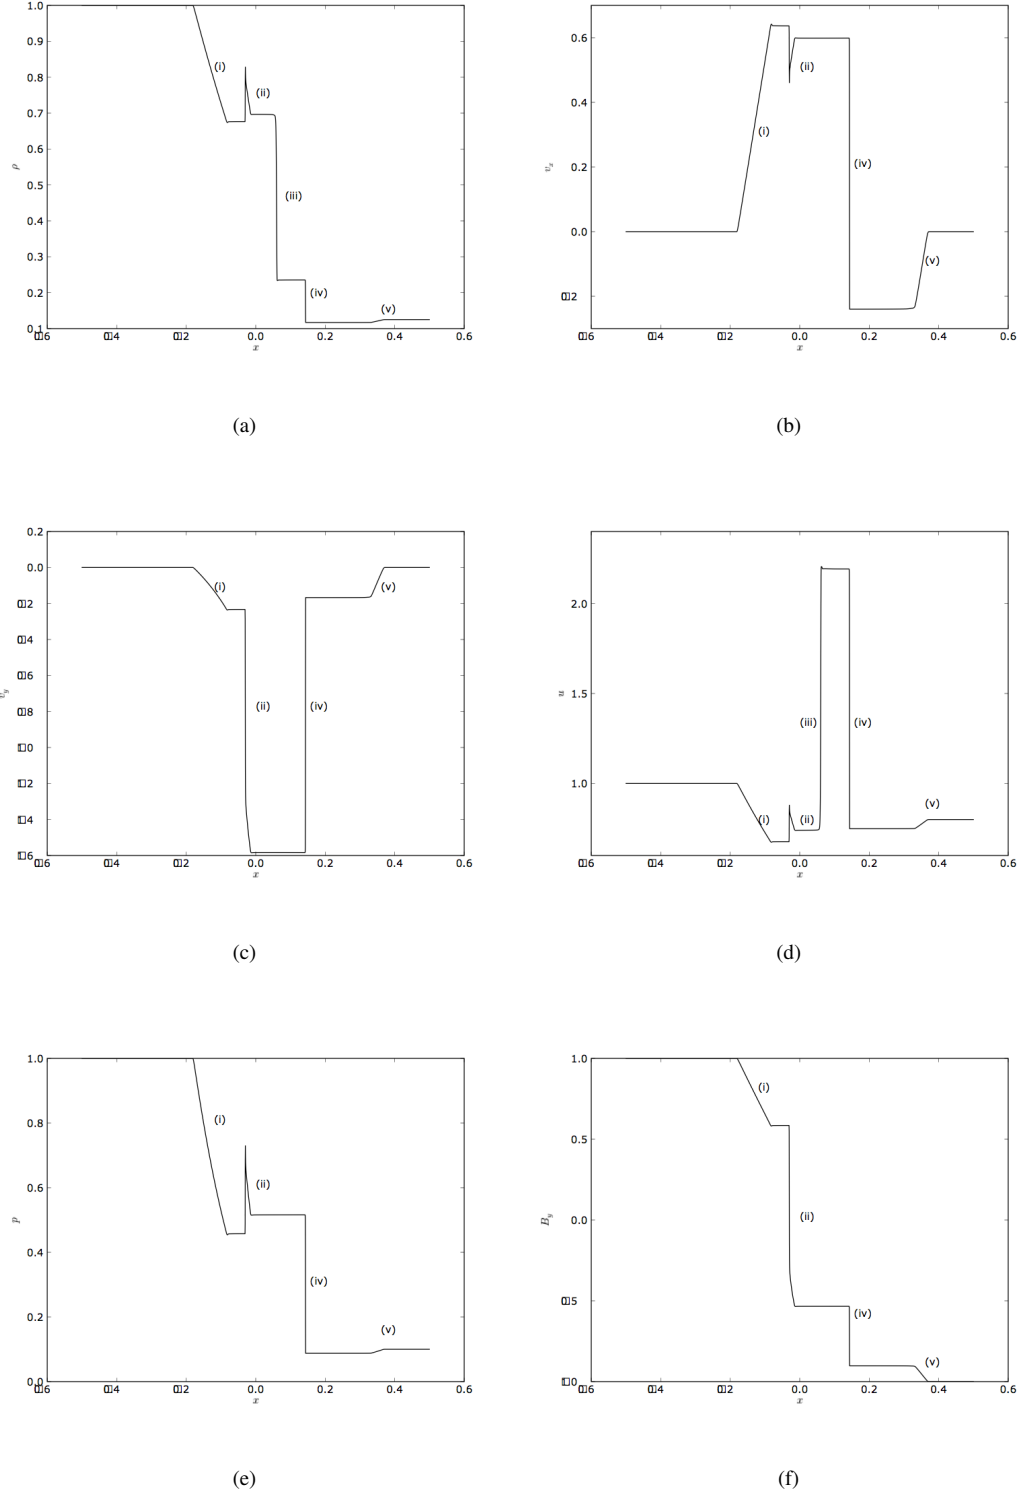


Figure 2.1: The waves present in (a) ρ , (b) v_x , (c) v_y , (d) u , (e) p , and (f) B_y in an MHD shock tube: (i) left-moving fast rarefaction, (ii) left-moving slow compound wave (shock/rarefaction), (iii) contact discontinuity, (iv) right-moving slow shock, and (v) right-moving fast shock.

material itself and assuming that its topology is invariant. For this reason, it is often said that magnetic field lines are “frozen into” perfectly conducting media.

In a real conductor, the magnetic field is allowed to “slip through” the medium at a rate that increases with that medium’s electrical resistivity η . This phenomenon is qualitatively similar to thermal diffusion: regions of intense magnetic field contribute to the field in regions of less intensity until magnetostatic equilibrium is achieved. Any magnetic field energy lost in this process is converted to thermal energy by Joule heating [21]:

$$\frac{\partial u}{\partial t} = \mathbf{J} \cdot \mathbf{E} = \eta J^2 \quad (2.57)$$

If we define the magnetic “diffusivity” as $D_m = \max(\eta_{ii})/\mu_0$, then the characteristic time over which a magnetic field decays within a conductor with no external field can be estimated using the usual time scale for diffusion in a system of characteristic length L :

$$\tau \approx \frac{L^2}{D_m} \quad (2.58)$$

Magnetic reconnection

Because the magnetic field lines are not frozen into a real conductor, it is possible for the topology of the magnetic field to change. Thus, field lines that are not connected at a certain time may become connected at a later time, and those that are previously connected may disconnect. This magnetic reconnection is a topic of much interest in the field of plasma physics, since the orientations of different reconnecting magnetic field lines affect the forces felt by the plasmas that contain them.

2.5 Limitations of MHD

We have used several simplifying assumptions in pursuit of MHD, and now we must examine the limitations of the resulting equations above and beyond those of the fluid approximation. Firstly, MHD cannot

describe any phenomena that depend upon the detailed motion of electrons, such as cyclotron orbits, plasma frequency oscillations, electromagnetic waves, and the like. This is because we decided that the electrons can respond instantaneously to the magnetic field, ignoring the time scale at which this response occurs.

Secondly, MHD cannot describe plasmas for which kinetic effects are important, since it assumes that the distribution function for each species is Maxwellian, and therefore in or close to local thermal equilibrium.

Finally, we have only discussed neutral and quasi-neutral plasmas in the sense that there are assumed to be equal numbers of positive and negative charges in the plasma. Otherwise there would be a net electric field arising from the excess charge.

For a more detailed discussion of the systems most accurately described by MHD, see Hazeltine [61].

2.6 Requirements for Modeling Plasmoids

We conclude this chapter with a list of plasma features that must be accurately represented by any simulation of laboratory plasmoids. As we discussed in the introduction, the simulation of plasmas (even under ideal or simplified circumstances) is a challenging problem that remains to be solved satisfactorily. There is no numeric method that can perfectly capture plasma dynamics. Therefore it is crucial that we explicitly choose the aspects of our model that we deem necessary to reproduce.

1. The presence and speeds of linear MHD waves

Any credible simulation of an MHD plasma must reproduce the fundamental solutions that we have discussed: fast and slow compressional (magnetosonic) waves and longitudinal Alfvén waves. These waves should propagate linearly at their predicted speeds under the idealized circumstances in which they were derived.

2. Shocks and discontinuities

The currents and magnetic fields used to accelerate plasmas in the laboratory can generate signals that are much larger than the small signals we assumed for the linear regime of the fluid equations, so these signals

will tend to steepen into shocks. Even if the plasma's initial conditions are comfortably linear from a hydrodynamic perspective, strong currents will drive discontinuities in the magnetic field that will in turn generate strong driving $\mathbf{J} \times \mathbf{B}$ forces. We should observe the development of fast and slow shocks where these currents are strong enough.

3. Magnetic flux conservation and diffusion

In a perfect conductor, the magnetic flux should be conserved to sufficient numeric precision that the dynamics are not noticeably affected. In a resistor, the magnetic field must diffuse at the predicted rate, and the resistivity of the medium should be allowed to vary according to physical models such as the Spitzer model [70]. For us, this diffusion rate is much more important than absolute flux conservation, since real plasmas are not ideal conductors and since this rate determines the ratio of conversion of magnetic energy to kinetic/thermal energy. Ideal MHD is much more important to the astrophysics community for its role in large magnetic fields [71].

4. Physical boundaries

The geometry and material composition of the device that contains a plasma necessarily has a large effect on that plasma's shape and behavior. Rigid, impermeable bodies can cause large signals in fluids; conducting walls repel magnetic fields, whereas insulating walls allow these fields to escape. We must develop a realistic set of boundary conditions that represent the laboratory devices in which our plasmas exist.

5. The plasma-vacuum interface

Because laboratory plasmas are typically created in a vacuum environment, we must be able to represent the interface separating the plasma from the vacuum. It is along this interface that surface currents will develop, and these surface currents must be represented accurately, as they generate the all- important $\mathbf{J} \times \mathbf{B}$ force that accelerates the plasma.

6. The $\nabla \cdot \mathbf{B} = 0$ constraint

Toth [26] and others ([72], [73]) have discussed the importance of maintaining a divergence-free magnetic field. If this constraint is ignored, aphysical forces can arise that cause the behavior of a plasma to diverge from that of its real counterpart. Unfortunately, this constraint is not enforced by the MHD equations themselves. In the introduction, we discussed some of the ways in which different groups have addressed this issue. We must follow suit.

As we develop the SPH method that forms the bulk of the work in this thesis, we will describe test problems that we have used to verify each one of these requirements.

Chapter 3

Smoothed Particle Hydrodynamics

In this chapter we introduce the basic ideas underlying Smoothed Particle Hydrodynamics (SPH) and its application to problems in ordinary fluid dynamics. We describe how the fluid equations are transformed from partial differential equations to ordinary differential equations and then integrated in time. The art of shock capturing is crucial for any treatment of fluid mechanics, and we explain the challenges and methods most commonly used to address them in SPH. We then enumerate the issues and the numerical parameters that are most relevant to typical SPH simulations and list some refinements to the technique that can improve its accuracy. We conclude the chapter with a list of test problems relevant to simulations of compressible fluid flows.

3.1 The SPH Formulation

Like most mesh-free methods, SPH represents quantities in space with a set of N disconnected points in space called *nodes* that have overlapping regions of influence. The quantities are defined on the nodes themselves, and spatial derivatives are defined through an interpolation process involving a set of neighbors $\{j\}$ for each node i that fall within the region of influence Ω_i , as shown in Figure 3.1.

Thus, the connectivity (defined by set of nodes $\{j\}$ for each i) changes as those nodes pass through




Figure 3.1: The “stencil” of a node in SPH in 2D. The radius of influence is determined by the desired number of neighbors per smoothing scale, n_h : (a) $n_h = 1.2$, (b) $n_h = 1.4$, (c) $n_h = 1.8$, (d) $n_h = 2.0$.

different regions. One can easily draw parallels between this type of method and a Free-Lagrange reconnection method [74]. Methods of this sort tend to be more expensive than those with fixed connectivity but often handle problems with large deformations more gracefully.

Note that the “radius of influence” of a node above has nothing to do with physics. It is purely a measure of the extent of the finite derivatives computed by the SPH algorithm. All of the physical length parameters are accommodated by our choice of the fluid equations that we derived in the last chapter. The nodes we discuss are fluid elements, not particles, and the fluid equations cannot describe phenomena at length scales below those at which a material is continuous.

Our treatment of SPH here is brief; several thorough expositions exist in the literature ([75], [76], [77]). We begin with the Dirac delta function representation of a scalar function f [78] on a domain $x \in [x_1, x_2]$:

$$f(x) = \int_{x_1}^{x_2} f(x') \delta(x - x') dx' \quad (3.1)$$

In SPH we replace $\delta(x - x')$ with a *kernel function* $W(x - x'; h)$ where h is a parameter that characterizes the region of influence around a node.

$$f(x) \approx \int_{x_1}^{x_2} f(x') W(x - x'; h) dx' \quad (3.2)$$

As $h \rightarrow 0$, $W(x - x'; h) \rightarrow \delta(x - x')$ and this approximation gets better. We refer to h as the *smoothing length* or *smoothing scale*. A typical form of the smoothing kernel is the B-spline kernel, defined by ([76])

$$W(|\mathbf{x} - \mathbf{x}'|; h) = W(\eta; h) = A_v \begin{cases} 1 - \frac{3}{2}\eta^2 + \frac{3}{4}\eta^3, & \eta \leq 1 \\ \frac{1}{4}(2 - \eta)^3, & \eta \leq 2 \\ 0, & \eta > 2 \end{cases} \quad (3.3)$$

where $\eta \equiv |\mathbf{x} - \mathbf{x}'|/h$ and A_v is a normalization constant defined so that $\int W(|\mathbf{x} - \mathbf{x}'|; h) d\Omega = 1$ in v dimensions. The task remains to compute the integral (3.2). The usual approach in SPH is to break the region up into N elements, with the i th element having a length (or more generally, a volume) $\Delta V_i = \frac{m_i}{\rho_i}$, where m_i is the

mass of the element and ρ_i is its mass density (mass per unit volume). The approximation of $f(x)$ is

$$f(x) \approx \sum_{i=1}^N \frac{m_i}{\rho_i} f(x_i) W(x - x_i; h) \equiv \langle f(x) \rangle \quad (3.4)$$

Note that our definition of a volume element $\Delta V_i = m_i / \rho_i$ gives us a natural expression for the mass density of the fluid. If we substitute ρ for f in (3.4), we obtain

$$\langle \rho(x) \rangle \equiv \sum_{i=1}^N m_i W(x - x_i; h) \quad (3.5)$$

We refer to this expression as the *mass density summation formula*. For multi-dimensional flows, the SPH approximation formula is

$$\langle f(\mathbf{x}) \rangle \equiv \sum_{i=1}^N \frac{m_i}{\rho_i} f(\mathbf{x}_i) W(|\mathbf{x} - \mathbf{x}_i|; h) \quad (3.6)$$

and the mass density summation formula is

$$\langle \rho(x) \rangle \equiv \sum_{i=1}^N m_i W(|\mathbf{x} - \mathbf{x}_i|; h). \quad (3.7)$$

In any case, the approximation of f at x is accurate to second order in Δx so long as f is smooth and defined on points that are relatively ordered, and so long as the kernel function W is symmetric about the origin. This has been shown by several authors ([76], [75]). It is instructive to keep in mind these stipulations when working with variants of SPH with asymmetric kernels or when simulating highly-disordered flows.

It is tempting to think of (3.4) as an interpolation method, but in fact it is straightforward to show that $f(x_i) \neq \langle f(x_i) \rangle$. At this point it is appropriate to discuss the *consistency* of an approximation. A method is said to be *consistent* with respect to a class of functions if it exactly reproduces that class of functions at the defined points. In particular, a method is *consistent to polynomial order p* if it exactly reproduces polynomial functions up and including those of p th order. Consistent methods are attractive because they do not introduce spurious changes in quiescent systems, and because their convergence rates depend less on the disorder of the

underlying points. These methods are typically also more expensive, though the extra computational work required varies greatly between approaches.

In this language, SPH is *inconsistent* in the sense that it does not exactly reproduce any class of functions: constant, linear, quadratic, or otherwise. The most unfortunate consequence of this inconsistency is a numerical instability that manifests itself when nodes are attracted to one another and causes them to clump together. This is called the *tensile instability* [79] and has been identified in the context of the consistency of spatial derivatives [80], Von Neumann stability analysis [81], and the shape of the smoothing kernel [82]. While its effects are not always evident in ordinary fluid flows, it is a source of major concern to anyone using SPH, particularly for elastodynamics [83] or for magnetohydrodynamics [84] since both of these applications involve appreciable attractive forces. There are several variants of SPH ([85], [86]) and other mesh-free methods ([87], [5]) that are designed to be consistent. We refer the interested reader to [77] for a discussion on the notion of consistency itself. Meanwhile, there are other approaches to addressing the tensile instability in the less exotic variants of SPH (e.g. [79]) with varying degrees of success.

3.1.1 Spatial discretization

The derivative of f with respect to x may be computed by differentiating (3.1) and integrating by parts. This allows the derivative to be expressed in terms of the derivative of the kernel function:

$$\frac{df(x)}{dx} = \int_{x_1}^{x_2} \frac{df}{dx'}(x')W(x-x')dx' \quad (3.8)$$

$$= fW \Big|_{x_1}^{x_2} - \int_{x_1}^{x_2} f(x') \frac{dW}{dx'}(x-x';h)dx' \quad (3.9)$$

$$\approx fW \Big|_{x_1}^{x_2} + \sum_{i=1}^N \frac{m_i}{\rho_i} f(x_i) \frac{dW}{dx'}(x-x_i;h) \quad (3.10)$$

Above, we have observed that $\frac{dW}{dx'}(x-x') = -\frac{dW}{dx}(x-x')$. It is common practice to neglect the surface term completely on the grounds that $W \rightarrow 0$ at the boundary. As we will see later, there are many point distributions for which this is simply not true. In these situations, we can either treat the boundary term explicitly or add “ghost nodes” outside the boundary that will mimic the properties of the boundary. We generally adopt the

latter approach.

We have derived the SPH approximation for one-dimensional geometry, but this same approach can be used for scalar, vector, and tensor-valued functions in one, two, and three dimensions. Thus one can derive the following (simple but suboptimal) expressions for the gradient of a scalar field and the vector derivatives of a vector field:

$$\langle \nabla f \rangle_i = \sum_j \frac{m_j}{\rho_j} f_j \nabla W_{ij}(h_i) \quad (3.11)$$

$$\langle \nabla \cdot \mathbf{F} \rangle_i = \sum_j \frac{m_j}{\rho_j} \mathbf{F}_j \cdot \nabla W_{ij}(h_i) \quad (3.12)$$

$$\langle \nabla \times \mathbf{F} \rangle_i = \sum_j \frac{m_j}{\rho_j} \mathbf{F}_j \times \nabla W_{ij}(h_i) \quad (3.13)$$

Here and hereafter we use some notation commonly found in the SPH literature (including surrounding the SPH approximation of a spatial derivative with angle brackets $\langle \rangle$):

$$f_i \equiv f(\mathbf{x}_i) \quad (3.14)$$

$$\mathbf{x}_{ij} \equiv \mathbf{x}_i - \mathbf{x}_j \quad (3.15)$$

$$\mathbf{v}_{ij} \equiv \mathbf{v}_i - \mathbf{v}_j \quad (3.16)$$

$$W_{ij}(h_i) \equiv W(\mathbf{x}_i - \mathbf{x}_j; h_i) \quad (3.17)$$

$$W_{ij}(h_j) \equiv W(\mathbf{x}_i - \mathbf{x}_j; h_j) \quad (3.18)$$

$$W_{ij} \equiv \frac{1}{2}(W_{ij}(h_i) + W_{ij}(h_j)) \quad (3.19)$$

Summations over neighbors j of a node i involve only those neighbors for which the corresponding kernel is nonzero. We have introduced a nodal smoothing scale h_i so that these neighbors j can be chosen dynamically depending on the density of the nodes. We will discuss the notion of a spatially-varying smoothing scale later in this section. Note that we have several expressions involving the smoothing kernel. These have to do with symmetric and unsymmetric forms of the kernel which are used in variants of SPH.

The expressions for the spatial derivatives are certainly not unique. Based on the specific needs of an algorithm, we can define different approximations for these spatial derivative operators. We describe some of the more common examples below:

- **Symmetric gradient**

The gradient ∇f of a scalar function is antisymmetric in the sense that $\frac{\partial f}{\partial x}(x - x') = -\frac{\partial f}{\partial x}(x' - x)$. This has important physical consequences. For example, a conservative force represented by a gradient, $\mathbf{F} = \nabla\Phi$ results in dynamic motion in which energy and momentum are conserved. The SPH gradient we have derived above does not have this property because

$$\frac{m_i}{\rho_i} \nabla W_{ij}(h_i) \neq -\frac{m_j}{\rho_j} \nabla W_{ji}(h_j) \quad (3.20)$$

in general. However, we can construct a gradient that has the antisymmetric property [76]:

$$(\nabla f)_i = \rho_i \sum_j \left(\frac{m_i}{\rho_i^2} f_i + \frac{m_j}{\rho_j^2} f_j \right) \nabla W_{ij} \quad (3.21)$$

As we will see in the next section, the symmetry of the gradient operator is often too important to overlook numerically when conservation is relevant.

- **Mass-density-weighted forms**

The mass density appears in the denominator of the approximations (3.11), (3.12), and (3.13). This is something of a nuisance because we have to know the mass density at each point j before we compute these derivatives, and the accuracy of the approximation can suffer when the density changes suddenly in space [76]. A simple fix is to write the derivatives in terms of those of the product of the density with the quantities. For example, we can express ∇f in terms of $\nabla(\rho f)$:

$$\nabla \rho f = \nabla \rho f + \rho \nabla f \quad (3.22)$$

The resulting equations for the spatial derivatives in lieu of (3.11) - (3.13) are

$$\langle \nabla f \rangle_i = \frac{1}{\rho_i} \sum_j m_j f_j \nabla W_{ij}(h_i) \quad (3.23)$$

$$\langle \nabla \cdot \mathbf{F} \rangle_i = \frac{1}{\rho_i} \sum_j m_j \mathbf{F}_j \cdot \nabla W_{ij}(h_i) \quad (3.24)$$

$$\langle \nabla \times \mathbf{F} \rangle_i = \frac{1}{\rho_i} \sum_j m_j \mathbf{F}_j \times \nabla W_{ij}(h_i) \quad (3.25)$$

Thus, the mass density is pulled outside the sum, making the expressions less sensitive to fluctuations in mass density.

- **Number-density-weighted forms**

The mass-density weighted forms have many nice properties, but they aren't perfect. In regions in which the masses m_i and m_j of interacting nodes are very different, those expressions become quite inaccurate. Another variant that does better under these circumstances is the number-density-weighted form, in which we differentiate the quantity of interest times its number density $n_i \equiv \rho_i/m_i$. This gives us the new expressions

$$\langle \nabla f \rangle_i = \frac{m_i}{\rho_i} \sum_j f_j \nabla W_{ij}(h_i) \quad (3.26)$$

$$\langle \nabla \cdot \mathbf{F} \rangle_i = \frac{m_i}{\rho_i} \sum_j \mathbf{F}_j \cdot \nabla W_{ij}(h_i) \quad (3.27)$$

$$\langle \nabla \times \mathbf{F} \rangle_i = \frac{m_i}{\rho_i} \sum_j \mathbf{F}_j \times \nabla W_{ij}(h_i) \quad (3.28)$$

Unfortunately, this is no generic replacement for the mass-density weighted form, since it gives worse results for equations in which conservation is important.

Many other variants are available. Colin et al [88] have written a good discussion about the different forms of the gradient and Laplacian operators that demonstrate the different degrees of accuracy in different node configurations.

We have written the above equations in vector form to emphasize that the SPH approximation generates equations which are *Galilean-invariant*. That is, they do not depend upon the use of a specific

coordinate system. This is important, since any dependence of a method on a given coordinate system can introduce coordinate-dependent artifacts in its solutions.

3.1.2 The fluid equations

Consider a non-conducting fluid in which $\mathbf{B} = \mathbf{0}$. The Lagrangian fluid equations ((2.45), (2.46), (2.47)) reduce to

$$\frac{D\rho}{Dt} = \rho \nabla \cdot \mathbf{v} \quad (3.29)$$

$$\frac{D\mathbf{v}}{Dt} = -\frac{\nabla p}{\rho} \quad (3.30)$$

$$\frac{Du}{Dt} = -\frac{p}{\rho} \nabla \cdot \mathbf{v} \quad (3.31)$$

These can be expressed as ordinary differential equations in the SPH formalism in a relatively straightforward way using various expressions for the spatial derivatives given above. One typical configuration is

$$\frac{d\rho_i}{dt} = \sum_j m_j \mathbf{v}_{ij} \cdot \nabla W_{ij} \quad (3.32)$$

$$\frac{d\mathbf{v}_i}{dt} = \sum_j m_j \left(\frac{p_i}{\rho_i^2} + \frac{p_j}{\rho_j^2} \right) \nabla W_{ij} \quad (3.33)$$

$$\frac{du_i}{dt} = \frac{p_i}{\rho_i^2} \sum_j m_j \mathbf{v}_{ij} \cdot \nabla W_{ij} \quad (3.34)$$

Here we have used the usual form of the continuity equation, a symmetric- gradient version of the momentum equation, and an unsymmetric form of the energy equation. This may seem arbitrary, and it is certainly neither the best nor the worst set of equations to use, but there are good reasons for choosing these forms. For instance, the symmetry of the momentum exchange between nodes i and j in (3.33) guarantees that the total linear momentum of the system $\mathbf{P} = \sum_i m_i \mathbf{v}_i$ is conserved, since whatever is lost by node i is gained by node j . However, total energy $E = \sum_i m_i u_i$ is not conserved because there is no such symmetry in (3.34). A symmetric form of the energy equation is

$$\frac{du_i}{dt} = \frac{1}{2} \sum_j m_j \left(\frac{p_i}{\rho_i^2} + \frac{p_j}{\rho_j^2} \right) \mathbf{v}_{ij} \cdot \nabla W_{ij} \quad (3.35)$$

The problem with using (3.35) is that in the process of improving energy conservation, sometimes a cool node i will surrender some of its internal energy to a hotter node j . This is bad thermodynamics, and its adverse effects easily overshadow any gains from total energy conservation. There are other ways of improving energy conservation that can avoid this “artificial cooling,” but this requires more work than simply choosing a different SPH div operator.

Again, the specific form of the fluid equations will be dictated by the specific needs of a simulation. This approximation conserves linear momentum, making it a good candidate for problems involving shock waves. But if an application requires very precise energy conservation, this may not be the best set of SPH fluid equations to use.

3.1.3 The SPH smoothing scale

We have written the SPH spatial derivatives in terms of varying nodal masses $\{m_i\}$ and smoothing scales $\{h_i\}$. Early forms of SPH simply used a single value of h for all of the nodes, but this has been found to be unsatisfactory for flows in which the spacing of nodes vary significantly. The need for a dynamic and inhomogeneous value of h can best be understood from its relationship to the mass density of the node distribution as given by (3.7):

$$\rho_i = \sum_j m_j W_{ij}(h_i) \quad (3.36)$$

(Here, $\{j\}$ includes i itself.) This is in general a nonlinear relationship between ρ and h , and the best way to express h was not obvious to authors in the early days of SPH. One method was to define an evolution equation for h [76] by adopting the functional form

$$h_i \propto \left(\frac{m_i}{\rho_i} \right)^{1/\nu} \quad (3.37)$$

(where v is the number of dynamically significant spatial dimensions) and taking its time derivative:

$$\frac{dh_i}{dt} = m_i^{1/v} \frac{d}{dt} (\rho^{-1/v}) \equiv -\frac{h_i}{\rho_i} \frac{d\rho_i}{dt} \quad (3.38)$$

This technique is conceptually simple and computationally inexpensive (since $\frac{d\rho_i}{dt}$ can be computed at the same time) and is still in use today. But over time, the discretization error for h can grow to be very large as the nodal disorder increases.

More recently, authors have been focusing on the nonlinear relationship between ρ and h . One consistent definition for h is

$$h_i = n_h \left(\frac{m_i}{\rho_i} \right)^{1/v} \quad (3.39)$$

Here we have introduced a parameter, n_h , that dictates how many neighbors should fall within a single smoothing scale of a node. Thacker [89] and Owen use an iterative scheme to bring h into line with (3.39) based on the SPH kernel contributions of neighboring nodes and then use (3.7) to update the mass density so that it is consistent with h . Meanwhile, Monaghan and Price [90] solve (3.39) and (3.7) iteratively for h and ρ until some convergence threshold is reached. Both of these methods seem to work well for fluids and result in a mass density that is consistent with the node distribution. In this dissertation, we use both Monaghan's and Owen's approaches—their results do not differ all that much. Owen's method has the advantage that it can be extended easily to treat anisotropic smoothing scales, which we discuss briefly below.

The choice of n_h remains somewhat of a controversy in the SPH community. Most authors seem to run simulations with n_h at around 1 or 1.2. This results in faster runtimes, but it is not clear to us that the accuracy of such a value is sufficient. Our work with Owen suggests that higher values of n_h produce better convergence rates in general, and we are inclined to set n_h at around 2 for hydrodynamics problems. On the other hand, Martin et al [82] do calculations that suggest that values of n_h higher than 1.4 are theoretically unstable, but we have not found this practically to be the case in hydrodynamic flows.

We have indeed found that some MHD problems exhibit instability at $n_h = 2$, but these instabilities

are problem-dependent and their convergence rates still improve with higher values of the parameter. Price seems to share concerns regarding n_h and stability with the other authors [84], and sets n_h to around 1.2 to 1.4 for most of his simulations. Perplexingly, there are a few cases in which he chooses values that exceed even ours (such as in [3], where $n_h = 2.4$ is used) with no explanation.

Clearly, there is more work to be done in exploring this parameter and with the smoothing length in general. We describe our own explorations in the section on test problems.

3.2 Shock Waves

The momentum equation (2.40) is nonlinear in \mathbf{v} , giving rise to a number of phenomena that are absent in purely linear systems. If we consider very weak sound waves, we can linearize the equation to obtain acoustic wave solutions, which are continuous and which obey the principle of superposition, placing hydrodynamics into a well-behaved mathematical context familiar to practitioners of linear electrodynamics. However, sound waves of measurable amplitude will eventually interact with one another via the nonlinear term $\mathbf{v} \cdot \nabla \mathbf{v}$ and generate new sound waves of higher frequency [91]. This phenomenon is called *nonlinear harmonic generation*, and it has the effect of “steepening” sound waves until they become discontinuous. Figure 3.2 shows an example of this steepening.

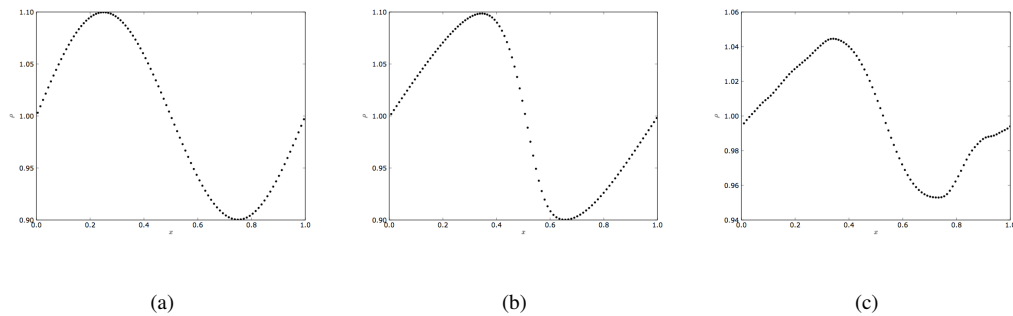


Figure 3.2: The gradual steepening of a linear wave. (a) is the original acoustic wave, and (b) and (c) show the distortion of the wave through the $\mathbf{v} \cdot \nabla \mathbf{v}$ term.

There are two types of discontinuities in hydrodynamic flows. The first type is a *contact disconti-*

nuity: an interface separating media of two different densities at which the pressure is balanced. There is no flow through such an interface. The second type is a *shock wave*, which is a discontinuity in pressure across which mass flows. The statements we make about treating numerical discontinuities apply to both types, but we will focus on treating shocks in this section.

A physical way to understand the development of a shock wave is to think of the sound speed of a medium $c_s = \sqrt{P/\rho}$ as the speed limit of all signals in that medium. Then any energetic event (the release of a barrier between two media far from equilibrium, or an explosion) can only propagate energy through a medium as fast as its sound speed. The signals carrying energy “pile up” against the sound speed barrier as they hit it, creating a discontinuity.

In the real world, the pressure and density of a fluid are not really discontinuous, since the molecules at a shock front are always in contact with one another, and any dissipative forces between them will tend to smooth the signal across the interface. However, the thickness of a shock is typically well below the length scale at which the fluid approximation is valid, so we must contend with how to represent these discontinuities.

3.2.1 The Rankine-Hugoniot Relations

The spatial derivatives of the mass density, the velocity, and the specific thermal energy cease to exist at a discontinuity, so we cannot rely on the the fluid equations (2.45) - (2.47) to characterize their solution there.

The equations describing the conservation of mass, momentum, and energy near these discontinuities are the *Rankine-Hugoniot relations* [66]:

$$\rho_1 v_1 = \rho_2 v_2 \quad (3.40)$$

$$p_1 + \rho_1 v_1^2 = p_2 + \rho_2 v_2^2 \quad (3.41)$$

$$u_1 + \frac{p_1}{\rho_1} + \frac{v_1^2}{2} = u_2 + \frac{p_2}{\rho_2} + \frac{v_2^2}{2} \quad (3.42)$$

where the subscripts 1 and 2 indicate quantities on either side of a discontinuous interface and v is the component of velocity incident upon the interface. These equations can be obtained by integrating the fluid equations (2.39) - (2.41) over a small region $[x_1, x_2]$ containing the discontinuity and taking the limit as $x_1 \rightarrow x_2$. For an ideal gas with specific heat ratio γ in which $p = \rho(\gamma - 1)u$, the velocities can be eliminated from the equations to obtain a relationship between the pressures:

$$\frac{p_1}{p_2} = \frac{(\gamma + 1) - (\gamma - 1)\frac{\rho_2}{\rho_1}}{(\gamma + 1)\frac{\rho_2}{\rho_1} - (\gamma - 1)} \quad (3.43)$$

The requirement that the pressures be positive means that the ratio of densities will never exceed $(\gamma + 1)/(\gamma - 1)$.

We have mentioned that dissipative processes in molecular physics determine the nature of shocks in the real world. To see how these processes enter the Rankine-Hugoniot relations, we consider what happens to the entropy, or disorder, of the fluid as it crosses the shock front, following the treatment given by Zeldovich [66]. The entropy of an ideal gas is defined, to within a constant, as

$$S = c_v \ln pV^\gamma \quad (3.44)$$

where c_v is the specific heat of the gas at constant volume and V is its *specific volume* or volume per unit mass, $V = \rho^{-1}$. We can use (3.43) to obtain a relationship between V_1 and V_2 :

$$\frac{V_1}{V_2} = \frac{(\gamma - 1)p_1 + (\gamma + 1)p_0}{(\gamma + 1)p_1 + (\gamma - 1)p_0} \quad (3.45)$$

Then the change in density crossing the shock front is

$$S_2 - S_1 = c_v \ln \frac{p_2 V_2^\gamma}{p_1 V_1^\gamma} \quad (3.46)$$

$$= c_v \ln \left(\frac{p_2}{p_1} \left[\frac{(\gamma - 1)(p_2/p_1) + (\gamma + 1)}{(\gamma + 1)(p_2/p_1) + (\gamma - 1)} \right]^\gamma \right) \quad (3.47)$$

As p_2/p_1 grows, so does the jump in entropy across the surface. Thus as fluid passes through the shock, its kinetic energy is converted to thermal energy. In order that the Rankine-Hugoniot relations be satisfied, there

must be some dissipative process present to perform this conversion. The details of this process determine only the width across which the dissipation occurs—its existence alone is enough to satisfy the conservation equations. In reality, the dissipation results from collisions between molecules or ions, and the distance over which the shock is spread is comparable to the mean free path of these particles. However, this scale length is far below that achievable in numerical hydrodynamics.

The behavior of fluids at shock fronts is considered so important that a large literature has developed around the *Riemann problem*, in which two fluids in different states are placed in contact with one another and allowed to evolve shocks. Accordingly, an entire class of numerical methods called *Riemann solvers* has been developed based on the Rankine-Hugoniot relations. In these solvers, a grid contains a piecewise-constant collection of “Riemann states” joined together that satisfy the relevant equations in their analytic form and evolve the boundaries of these states in time according to their signal velocities. We do not attempt to describe these methods in any detail here—we only pause to note the following about Riemann solvers:

- They treat one-dimensional shocks so well that they are often used as benchmarks for problems that have no known solutions. In fact, we check some of our 1D MHD test results against those obtained by a Riemann solver in the next chapter.
- They have been used in multi-dimensional flows with some success in spite of the fact that no analytic solution exists to the Riemann problem in 2D or 3D.
- Since they are based on analytic solutions of the R-H relations, they are dependent on the specific form of the equations of state of the media involved.
- Their equations are not Galilean-invariant: if the shock is not aligned with the coordinates in which the Riemann problem is solved, errors will be introduced in the flux terms.

Two examples of techniques based on the Riemann problem are Godunov’s method [27], which uses the exact solution to the Riemann problem to express fluxes between regions, and Roe’s method [92], which obtains these fluxes by solving an approximate eigenvalue problem. We will use a code based on the former to verify some of our hydromagnetic test problems in Chapter 4.

3.2.2 Artificial Viscosity

We have seen that dissipative mechanisms must be present for conservation of mass, momentum, and energy in the presence of shocks. The fluid equations we have discussed so far do not have any dissipative terms. In the simplest sense, shock capturing is achieved by adding terms to $\frac{d\mathbf{v}}{dt}$ and $\frac{du}{dt}$ that smooth \mathbf{v} and u across the shock and allow kinetic energy to be converted to thermal energy. If we fail to do this when we solve the fluid equations, oscillations will develop in the region of the fluid through which the shock has passed, called the *post-shock* region. These oscillations can be understood physically as an excess of kinetic energy, and mathematically as the Gibbs phenomenon that occurs when a function cannot be represented as a partial sum [78]. In SPH, the addition of a pairwise artificial viscosity Π_{ij} yields the following momentum and energy equations:

$$\frac{d\mathbf{v}_i}{dt} = \sum_j m_j \left(\frac{p_i}{\rho_i^2} + \frac{p_j}{\rho_j^2} + \Pi_{ij} \right) \nabla W_{ij} \quad (3.48)$$

$$\frac{du_i}{dt} = \sum_j m_j \left(\frac{p_i}{\rho_i^2} + \Pi_{ij} \right) \mathbf{v}_{ij} \cdot \nabla W_{ij} \quad (3.49)$$

where $\Pi_{ij} = (\Pi_i + \Pi_j)/2$ so that momentum is conserved. The most popular form of Π_i was introduced by Gingold and Monaghan [93] and is based on the bulk and Von Neumann-Richtmyer viscosities:

$$\rho_i \Pi_i = \alpha c_{si} \mu_i + \beta \mu_i^2, \quad \mu_i = \frac{\mathbf{v}_{ij} \cdot \boldsymbol{\eta}_i}{|\boldsymbol{\eta}_i|^2 + \epsilon^2} \quad (3.50)$$

where $\boldsymbol{\eta}_i$ is the normalized node displacement $(\mathbf{x}_i - \mathbf{x}_j)/h_i$ and ϵ is a small parameter around 0.01. This model captures shocks admirably in 1D, confining the smoothing to within $2h$ of the shock front while providing sufficient resolution for the discontinuity itself to be represented. We shall see this in the test problems at the end of the chapter.

Most of the difficulty in applying artificial viscosity involves deciding where to do the smoothing. At its crudest, a model will drastically over-smooth the quantities of interest, greatly reducing the accuracy of the solution. This smoothing can be observed even in purely linear flows like acoustic waves or linear velocity profiles.

Another case of this oversmoothing occurs in the case of shear flow, in which $\nabla \cdot \mathbf{v} = 0$. Since the flow is not compressional, the “right thing” is for the artificial viscosity to contribute nothing to the fluid equations. Nevertheless, the unmodified Monaghan-Gingold viscosity is nonzero in such shear regions because the nodes do approach one another. Miller [?] and later Balsara [81] developed a simple correction factor to that is 1 under pure compression and 0 under pure shear:

$$f_i = \frac{\langle \nabla \cdot \mathbf{v} \rangle_i}{|\langle \nabla \cdot \mathbf{v} \rangle_i + \langle \nabla \times \mathbf{v} \rangle_i| + \epsilon^2 c_{si}/h_i} \quad (3.51)$$

More recently, Owen [94] has developed a tensor-valued model for artificial viscosity that further improves the ability of the model to distinguish between shear and compression.

Morris [95] reduces the artificial smoothing in hydrodynamic flows by time-evolving the viscosity parameter separately at each node, using source terms involving measurements of the compression of the fluid. He varies the linear viscosity parameter α_i at each node i , keeping the nonlinear parameter β constant across all nodes. α is time-evolved according to

$$\frac{d\alpha_i}{dt} = -\frac{\alpha_i - \alpha_{min}}{\tau_i} + (2 - \alpha_i) \max(-\nabla \cdot \mathbf{v}, 0) \quad (3.52)$$

and integrated in lockstep with all other time-integrated quantities as explained in the next section. The first term on the right-hand side is a “sink” term that attenuates the smoothing parameter in the absence of compression, and the second is a compressive “source” term.

Our results with Price’s choice of parameters indicate that varying these parameters attenuates the smoothing too much, leaving significant oscillations in the results of many tests. We will demonstrate these oscillations in the test problems. Price acknowledges this in his thesis [84]; his solution is simply to peg the parameters at a sufficiently high value to provide enough smoothing. Of course, this produces results that are oversmoothed.

Evidently, there remains much work to be done in the development of artificial viscosity models in SPH. In particular, a better limiter is needed to decide how much each region of the problem needs to be

smoothed.

On the Difference Between Artificial Viscosity and Riemann Solvers

Some numerical practitioners find offensive the introduction of ad-hoc dissipation terms in the fluid equations to “smooth over” the regions of difficulty. To address this scepticism, some authors have produced comparisons of SPH forms of artificial viscosities to the terms that appear in Riemann solvers([96], [97]). In this framework, nodes on either side of a discontinuity are analogous to Riemann states, and a dissipation term is constructed to account for the jumps in the given variables. The proposed SPH form is an extra term in the acceleration of a node due to dissipation.

$$\left(\frac{d\mathbf{v}_i}{dt}\right)_{visc} = \sum_j m_j \frac{\alpha v_{sig} \mathbf{v}_{ij} \cdot \hat{\mathbf{r}}_{ij}}{2\bar{\rho}_{ij}} \nabla W_{ij}(h_i) \quad (3.53)$$

accompanied by an extra heating term in the energy equation

$$\left(\frac{du_i}{dt}\right)_{visc} = - \sum_j m_j \frac{\alpha v_{sig} [\mathbf{v}_{ij} \cdot \hat{\mathbf{r}}_{ij}]^2}{4\bar{\rho}_{ij}} \hat{\mathbf{r}}_{ij} \cdot \nabla W_{ij}(h_i) \quad (3.54)$$

Above, $\hat{\mathbf{r}}_{ij} = (\mathbf{x}_i - \mathbf{x}_j)/|\mathbf{x}_i - \mathbf{x}_j|$, $\bar{\rho}_{ij} = (\rho_i + \rho_j)/2$, and v_{sig} is a signal velocity describing the relative velocity of approach of two nodes i and j , e.g.:

$$v_{sig} = c_{si} + c_{sj} - \beta \mathbf{v}_{ij} \cdot \hat{\mathbf{r}}_{ij} \quad (3.55)$$

This appears somewhat different from the Gingold-Monaghan form but it contains essentially the same linear and nonlinear coefficients α and β , and is the same order of magnitude. The two SPH forms give comparable results in the test problems. We extend this form to handle discontinuities in magnetic fields in Chapter 4.

Other Artificial Dissipation

Price [98] extends the usual artificial viscosity model to include artificial heat conductivity, applying variable parameters in a manner similar to those described above. He has argued in more recent work [99]

(in response to allegations that SPH cannot handle contact discontinuities [100]) that such heat dissipation is needed to treat discontinuities that appear in Kelvin-Helmholtz-type instabilities. The contribution of node j to the specific thermal energy at node i from artificial heat conduction is

$$\frac{\partial u}{\partial t}_{ij}^Q = \frac{1}{4} m_i m_j v_{sig} (\hat{\mathbf{r}} \cdot \nabla W_{ij}) \alpha_u (u_i - u_j) \quad (3.56)$$

For MHD problems, Price describes a similar mechanism he dubs “artificial resistivity” to smooth discontinuities in magnetic fields. We describe this in the next chapter.

3.3 Time Integration

(3.32) - (3.34) represent the *semi-discrete* system of fluid equations achieved by applying an SPH approximation to the continuous fluid equations. We are now free to choose the method by which we evolve ρ , \mathbf{v} , and u in time.

We divide a continuous time interval $[t_0, t_f]$ into a number of non-uniform segments, or *time steps*, $\{\Delta t_j\}$, where $\Delta t_j = t_{j+1} - t_j$. A quantity f at a node i computed at a time $t_n = t_0 + \sum_j^n \Delta t_j$ is written f_i^n . Likewise, the time derivative of f at node i evaluated at time t_n is $(df/dt)_i^n$. The time-indexed superscript can be used to distinguish between different time centerings in higher order time integration schemes.

We can demonstrate the time evolution of a quantity f using a few of the simplest time integration methods. The Forward Euler Method is the most straightforward way to compute f_i^{n+1} from f_i^n and $(df/dt)_i^n$:

$$\left(\frac{df}{dt}\right)_i^n = F(\{f_j^n\}, t_n) \quad (3.57)$$

$$f_i^{n+1} = f_i^n + \Delta t_n \left(\frac{df}{dt}\right)_i^n \quad (3.58)$$

This method is first-order accurate in Δt and is easy to implement [101]. It is also a *single-stage* integrator because time derivatives need only be computed at a single time (t_n) during a time step—this makes it easier to use for debugging algorithms. However, SPH is second-order accurate in space, so this method of integration

is too coarse for our purposes. Further, we have seen this method develop some numerical instabilities when run for long times on some problems. This is not surprising, since the time derivatives are computed at the beginning of the time step, analogous to a left hand quadrature rule. A midpoint-like rule, with df/dt evaluated at $t_{n+1/2}$, yields better results.

The most commonly-used methods for time-integrating the SPH equations are the second-order Runge-Kutta and Predictor-Corrector methods [101]. Both of these methods estimate the value of f at $t_{n+1/2}$ using a Forward- Euler-like method and then correct the calculation by using the the half-step data to compute f at t_{n+1} . The simulation results from these two methods are very similar to one another and are more accurate and stable than those of the Forward Euler method. However, both of these methods require two evaluations of the derivatives per time step, and are hence *multi-stage* integrators. Since the bulk of the work in SPH is performed in the calculation of time derivatives, these two integrators are effectively twice as expensive as single-stage integrators.

There is a price to be paid for using multi-stage integrators beyond the additional expense of computing more derivatives: in a single-stage integrator we may assume that the connectivity of the nodes remains constant over the span of a single time step. This can be a great advantage to certain algorithms, as we will see when we discuss refinements to the energy equation. For this reason, we have adopted a Leapfrog-like integration scheme that is a second-order single-stage integrator. In this scheme, f is centered at integer time steps n and the time derivative df/dt is evaluated only at each half-step $n + 1/2$:

$$f_i^{n+1/2} = f_i^n + \left(\frac{df}{dt}\right)_i^{n-1/2} \frac{\Delta t}{2} \quad (3.59)$$

$$\left(\frac{df}{dt}\right)_i^{n+1/2} = F(\{f_j^{n+1/2}\}, t_{n+1/2}) \quad (3.60)$$

$$f_i^{n+1} = f_i^n + \left(\frac{df}{dt}\right)_i^{n+1/2} \Delta t \quad (3.61)$$

The derivatives $\{df/dt_i^{n-1/2}\}$ are stored from the previous time step. At the first time step, $\{df/dt_i^{-1/2}\} \approx \{df/dt_i^0\}$. While less accurate at the same resolution as a second-order Runge-Kutta or Predictor-Corrector method, the “Leapfrog” method is much less expensive, performing neighbor searches and computing deriva-

tives only once per step. Further, it allows us to retain the assumption that node connectivity does not change over the course of a step.

In this chapter we use only *explicit* time integration methods, which are those for which f_i^{n+1} can be written completely in terms of quantities at time step n and those preceeding. *Implicit* methods (those for which f_i^{n+1} can depend on quantities at time step $n+1$ and those beyond) have been used before with SPH ([102], [103]), but the advantages that they confer do not seem to offset their cost for pure hydrodynamic flows. We will, however, use an implicit time integration scheme when we solve the magnetic diffusion equation in the next chapter.

3.3.1 Testing Convergence

We now demonstrate the convergence of the integration schemes we have discussed. It is important to establish this convergence if we are to assume that our single-stage Leapfrog integrator is able to replace the two-stage Runge-Kutta and/or Predictor-Corrector integrators. Because many of our hydrodynamic test problems include discontinuities or instabilities that prevent us from achieving optimal convergence, we test our implementations of these integrators using the linear ordinary differential equation

$$\frac{dy}{dt} - 2ty = t \tag{3.62}$$

$$y(0) = \frac{1}{2} \tag{3.63}$$

for which the solution is $y(t) = \exp(t^2) - 1/2$. The numeric solution \tilde{y} is evolved to $t = 1$ using a fixed number of steps and the global error $|\tilde{y} - y(1)|$ is computed. The global error should decrease with a power of the integration step size equal to its formal convergence rate. Table 3.1 gives the global error as a function of the number of steps N for several integrators, and Table 3.2 lists convergence rates of these global errors. We include a fourth-order Runge-Kutta algorithm in our study because it is the most commonly-used integrator for this kind of ODE.

The desired rates of convergence are achieved in this test. In particular, we see that the Leapfrog

N	FE	RK2	PC	LF	RK4
10	0.383648	0.019856	0.009225	0.041419	1.16×10^{-5}
20	0.207620	0.005306	0.002291	0.011434	7.45×10^{-7}
40	0.108314	0.001371	0.000570	0.002990	4.69×10^{-8}
80	0.055364	0.000348	0.000142	0.000763	2.94×10^{-9}
160	0.027995	8.77×10^{-5}	3.55×10^{-5}	0.000193	1.84×10^{-10}

Table 3.1: Global error $|\tilde{y} - y(1)|$ for FE (Forward Euler), RK2 (second-order Runge-Kutta), PC (Predictor-Corrector), LF (Leapfrog), and RK4 (fourth-order Runge-Kutta) integrators as a function of number of steps N

Integrator	Rate of convergence
FE	0.94600 ± 0.0018
RK2	1.95723 ± 0.0013
PC	$2.00583 \pm 8.5 \times 10^{-6}$
LF	1.93990 ± 0.0031
RK4	3.98833 ± 0.00017

Table 3.2: Convergence rates for global error $|\tilde{y} - y(1)|$ for FE (Forward Euler), RK2 (second-order Runge-Kutta), PC (Predictor-Corrector), LF (Leapfrog), and RK4 (fourth-order Runge-Kutta) integrators

algorithm has essentially the same convergence as its second-order Runge-Kutta and Predictor-Corrector counterparts, though the latter clearly outperforms the other two in accuracy for this test at any given resolution. Since it is the cheapest second-order algorithm (and the only single-stage one at that), we will use the Leapfrog integrator for all of our calculations in this thesis.

3.3.2 Constraints on Δt

The size of the n th time step, Δt_n , is chosen just before that time is taken, and is chosen in a way that guarantees both numerical stability and the successful capture of all interesting signals.

The CFL condition

Courant, Friedrichs, and Lewy [104] are credited with the insight that a signal moving at a speed c across a discrete interval Δx can only be represented accurately within a time interval Δt such that

$$\frac{\Delta x}{\Delta t} = c \quad (3.64)$$

This relation is called the Courant-Fridrichs-Lewy (or CFL) condition. In SPH the interval Δx between a node i and its neighboring nodes is approximately h/n_h . Since we are interested in sound waves, the propagation speed c is the sound speed c_s of the material, and

$$\Delta t \leq \min_i \frac{h_i}{n_h c_{si}} \quad (3.65)$$

Total velocity

The velocity itself is a limit on the time step in the sense that $|\mathbf{v}| \approx \Delta x / \Delta t$. Thus,

$$\Delta t \leq \min_i \frac{h_i}{n_h |\mathbf{v}_i|} \quad (3.66)$$

This prevents any node from traversing a distance greater than a nodal spacing within a single time step.

Velocity divergence

In a compressible medium, the divergence of the velocity is itself the inverse of the characteristic time of the expansion (if $\nabla \cdot \mathbf{v} > 0$) or collapse (if $\nabla \cdot \mathbf{v} < 0$) of that medium. This characteristic time provides another convenient upper limit for our time step:

$$\Delta t \leq \min_i \frac{1}{\langle \nabla \cdot \mathbf{v} \rangle_i} \quad (3.67)$$

where $\langle \nabla \cdot \mathbf{v} \rangle_i$ is the desired SPH interpolant of the velocity divergence.

Artificial viscosity

We have discussed some of the machinery necessary to capture shocks. Typically, the nodal spacing near shocks is significantly smaller than that elsewhere. In the light of the CFL condition, which limits the time step according to this nodal spacing and the sound speed, we imagine that regions with shocks will limit the time step more severely than others. We follow the common practice of using the artificial viscous

pressure Q to estimate a “viscous” sound speed $\sqrt{Q/\rho}$ with which we impose an analogous CFL condition for shock waves:

$$\Delta t \leq \min_i \frac{h_i}{n_h} \sqrt{\frac{\rho_i}{\Pi_i}} \quad (3.68)$$

Total acceleration

Unusually large accelerations can produce large changes in \mathbf{v} that thwart the constraints we have discussed so far, and so we temper these accelerations by observing that $|\mathbf{a}| \approx \Delta x / (\Delta t)^2$ so that time step is no more than

$$\Delta t \leq \min_i \sqrt{\frac{h}{n_h |\mathbf{a}_i|}} \quad (3.69)$$

Safety factor

In principle, a time step size that incorporates all of the above constraints will yield stable and accurate integration. In practice, a “safety factor” (between 0 and 1 and typically around 0.5) is commonly applied to the time step in order to provide a buffer of stability.

3.4 Issues and Refinements

At this point, we have described a basic SPH algorithm for solving the fluid equations. There remains much room for improvement to this basic scheme, and here we discuss ways in which various authors have addressed the most important issues.

3.4.1 Node Disorder

One of the difficulties of using mesh-free methods is that it is hard to characterize how well they perform when nodes become very disordered. While various mesh-free methods have been developed in an attempt to guarantee that the solutions remain robust under disordered nodal configurations ([50], [5]), it is

easier for us to address this issue by maintaining a level of order among the nodes. Monaghan does this by removing the constraint that the nodes move strictly with the fluid velocity [105]. In this method, an “XSPH” parameter χ between 0 and 1 is used to smooth the node velocities:

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{v}_i - \chi \frac{\sum_j (m_j/\rho_j) W_{ij}(h_i) \mathbf{v}_{ij}}{\sum_j (m_j/\rho_j) W_{ij}(h_i)} \quad (3.70)$$

XSPH is an extremely simple modification and helps to reduce node disorder fairly effectively. We employ it regularly in the test problems.

3.4.2 Stability

At the beginning of this chapter we discussed the notion of the consistency of a numeric approximation, and we noted that SPH was not consistent in that it cannot exactly reproduce any class of functions on a defined set of points. The biggest practical setback of this inconsistency is a “tensile” instability that occurs between nodes under the influence of an attractive force. There are many explanations for the tensile instability in the SPH literature ([79], [82], [84]), but the essential idea is that an SPH kernel function W is unable to keep nodes apart once they become sufficiently close to one another. This results in a “clumping behavior” that has been observed most severely in studies involving elastic materials and magnetic fields, both of which have tensile forces. The instability appears much less pronounced in regular hydrodynamic flows.

As of this thesis there is no universally-accepted technique for fighting this instability. Monaghan [79] meets the problem head-on, adding a repulsive pressure term to the momentum equation that gives clumping nodes a “kick” when they get too close to one another:

$$\frac{p_i}{\rho_i^2} + \frac{p_j}{\rho_j^2} + \Pi_{ij} \rightarrow \frac{p_i}{\rho_i^2} + \frac{p_j}{\rho_j^2} + \Pi_{ij} + R_i f_{ij}^n \quad (3.71)$$

where R is a pressure-dependent function designed to match the other pressure contributions at close range (e.g. $R_i \approx 0.2|p_i|/\rho_i^2$), $f_{ij} = W_{ij}(h_i)/W(1/n_h, h_i)$, and n is a small positive integer (e.g. 3, 4, 5). While

simple, this method seems to degrade the convergence of the solution, as we will show in the test problems. In his thesis, Price also explores the stability of this method and finds that while this term does prevent the unwanted clumping between nodes, the shape of the SPH kernel is altered significantly outside of the desired range [84].

A more thorough approach to addressing the tensile instability would be to use the Corrected SPH (CSPH) approximation of Bonet [80], in which additional “correction terms” are added to the expressions for the SPH kernel and its gradient. These terms add zeroth- and first-order consistency to SPH and place the technique on solid mathematical ground with other mesh-free methods used by the engineering community ([5], [85], [50]). However, some the SPH community have expressed concern that this adds too much computational expense.

The alternatives to these general methods for treating the tensile instability are specialized to treat the tensile forces that invite the instability. We will explore some treatments for dealing with magnetic tension in the next chapter, where the tensile instability plays a decidedly more major role.

3.4.3 Improved energy conservation

Conservation plays a large role in the formulation of the fluid equations, and the perfect numeric treatment of fluids would conserve mass, momentum, and energy to machine precision. The variant of SPH we have described conserves both mass and momentum in this way. The energy equation that we have described does not conserve energy, however. We are left with two options:

- conserve energy as well as we can within the discretization
- use another method to evolve the energy that perfectly conserves energy *and* yields non-negative nodal thermal energies.

Both options have been pursued in the SPH community. The first approach adds correction factors to the fluid equations that account for a spatially varying smoothing scale $h(\mathbf{x}, t)$. The second is a replacement

for the energy equation that expresses the transfer of energy between kinetic and thermal form, manifestly conserving energy.

∇h terms

It was noticed by Nelson and Papaloizou [106] and later by Hernquist and Springel [107] and Monaghan [108] that if h is allowed to vary in space, its gradient will contribute additional terms to the SPH equations through that of the SPH kernel. For example, Monaghan [108] notes that the SPH momentum equation may be written in Euler-Lagrange form as

$$\frac{d\mathbf{v}_i}{dt} = -\sum_j m_j \left(\frac{\partial u}{\partial \rho_s} \frac{\partial \rho}{\partial \mathbf{r}_a} \right) \quad (3.72)$$

where the s subscript refers to a quantity evaluated with the specific entropy s (entropy per unit mass) held constant. The density gradient at node j is

$$\frac{\partial \rho_j}{\partial \mathbf{r}_i} = \frac{\partial}{\partial \mathbf{r}_j} \left(\sum_k m_k W_{ik}(h_i) \delta_{ij} + m_i W_{ij}(h_j) \right) \quad (3.73)$$

and the kernel gradient, corrected for a spatially-varying h , is

$$\frac{\partial W_{ij}(h_i)}{\partial \mathbf{r}_i} = \nabla W_{ij}(h_i) + \frac{\partial W_{ij}(h_i)}{\partial h_i} \frac{\partial h_i}{\partial \mathbf{r}_i} \quad (3.74)$$

Instead of computing $\partial h / \partial \mathbf{r}$ directly, we can use the chain rule once more to relate it back to $\partial \rho / \partial \mathbf{r}$:

$$\frac{\partial h_i}{\partial \mathbf{r}_i} = \frac{\partial h_i}{\partial \rho_i} \frac{\partial \rho_i}{\partial \mathbf{r}_i} \quad (3.75)$$

and rewrite (3.73) in terms of $\nabla W_{ij}(h_i)$ and a correction factor Ω_i :

$$\frac{\partial \rho_j}{\partial \mathbf{r}_i} = \Omega_i^{-1} \sum_j \nabla W_{ij}(h_i) \quad (3.76)$$

where

$$\Omega_i \equiv 1 - \sum_j \frac{\partial W_{ij}(h_i)}{\partial h_i} \frac{\partial h_i}{\partial \rho_i} \quad (3.77)$$

and $\partial h_i / \partial \rho_i$ can be computed from the relation (3.39). The correction factor Ω_i is defined on each node and deviates from 1 only where effects from the gradient of h become important. Figure 3.3 illustrates the behavior of h and of Ω near a shock.




Figure 3.3: The behavior of (a) h , the SPH smoothing scale, and (b) Ω , the ∇h correction factor, near a shock.

Using (3.73) in (3.72) and a similar argument for the energy equation, the SPH fluid equations are modified to incorporate this correction:

$$\frac{d\mathbf{v}_i}{dt} = \sum_j m_j \left(\frac{p_i}{\Omega_i \rho_i^2} + \frac{p_j}{\Omega_j \rho_j^2} \right) \nabla W_{ij} \quad (3.78)$$

$$\frac{du_i}{dt} = \frac{p_i}{\Omega_i \rho_i^2} \sum_j m_j \mathbf{v}_{ij} \cdot \nabla W_{ij} \quad (3.79)$$

In Price's implementation, the SPH kernel is not symmetrized in the momentum equation as it is above.

Rather, each term on the right-hand side is multiplied by the kernel with the appropriate smoothing scale:

$$\frac{d\mathbf{v}_i}{dt} = \sum_j m_j \left(\frac{p_i}{\Omega_i \rho_i^2} \nabla W_{ij}(h_i) + \frac{p_j}{\Omega_j \rho_j^2} \nabla W_{ij}(h_j) \right) \quad (3.80)$$

We prefer in general to symmetrize the kernel where possible, since it gives results that are smoother.

In practice, we find that this correction does not make a significant difference in the accuracy of the results for pure hydrodynamic flows in 1D. This may be because we use a higher value of n_h for these problems. This correction becomes more much more significant in multidimensional flows, as we will see in the test problems. It may also add accuracy that compensates for the lower values of n_h that we are forced to use in magnetohydrodynamic flows.

A method for exact conservation

The energy equation is the expression of the conservation of energy in a continuum. As such, it governs the transfer of work into heat, and that of heat into work. Thus far, we have approximated this equation by replacing its continuous spatial derivatives with discrete SPH versions. Another approach is to find an expression that relates u_i^{i+1} to u_i^n in a way that exactly conserves energy and models the transfer of work and heat properly.

A numerical approximation of an equation that retains some important feature— be it the conservation of some quantity or the topological property of a derivative, say—is said to be *compatible* with or *mimetic* of that equation. Much fruitful work has been done in pursuit of compatible methods, in electromagnetics ([34], [109], [110]), where Maxwell’s equations rely on vector identities for the structure of their solutions, and in hydrodynamics ([111]) where conservation is important. Owen has recently developed an alternative to the SPH energy equation [112] that is compatible in the sense that it exactly conserves energy in hydrodynamic flows. For completeness, we give a very brief overview of his approach here.

The “compatible” energy formulation begins with the expression for the total energy in a hydrodynamic system at any given time n :

$$E^n = \sum_i m_i \left(\frac{1}{2} (v_i^n)^2 + u_i^n \right) \quad (3.81)$$

Thus, the change in total energy during a time step spanning $[t_{n+1}, t_n]$ which is zero if energy is conserved, is written

$$\Delta E^n = E^{n+1} - E^n = \sum_i m_i \left[\frac{1}{2} (v_i^{n+1})^2 + u_i^{n+1} - (v_i^n)^2 - u_i^n \right] = 0 \quad (3.82)$$

Introducing the acceleration $\mathbf{a}_i^n = d\mathbf{v}_i^n/dt$ and the time step size $\Delta t = t_{n+1} - t_n$ and expressing \mathbf{v}_i^{n+1} as $\mathbf{v}_i^n + \mathbf{a}_i^n \Delta t$, (3.82) becomes

$$0 = \sum_i m_i \left[\left(\mathbf{v}_i^n + \frac{1}{2} \mathbf{a}_i^n \Delta t \right) \cdot \mathbf{a}_i^n \Delta t + u_i^{n+1} - u_i^n \right] \quad (3.83)$$

$$\equiv \sum_i m_i \left[\mathbf{v}_i^{n+1/2} \cdot \mathbf{a}_i^n \Delta t + \Delta u_i \right] \quad (3.84)$$

where we have renamed the term in parenthesis $\mathbf{v}_i^{n+1/2}$, and defined $\Delta u_i = u_i^{n+1} - u_i^n$. This is a discrete statement of the conservation of energy with time centerings that guarantee that energy is conserved exactly. Finding the proper centerings of quantities in time and space is evidently key to formulating a compatible method, an observation first made by Yee [34].

The prescription is not complete—there are countless ways in which we could compute the changes $\{\Delta u_i\}$ so that (3.84) is satisfied. Owen writes the change in thermal energy in terms of the thermal interactions between pairs of neighboring nodes:

$$\Delta E^{thermal} = \sum_i m_i \Delta u_i = - \sum_i m_i \mathbf{v}_i^{n+1/2} \cdot \mathbf{a}_i^n \Delta t \quad (3.85)$$

$$= - \sum_i m_i \mathbf{v}_i^{n+1/2} \cdot \sum_j \mathbf{a}_{ij}^n \Delta t \quad (3.86)$$

where we have written $\mathbf{a}_i^n = \sum_j \mathbf{a}_{ij}^n$ in terms of the acceleration on node i from each neighbor j : \mathbf{a}_{ij} . If Δu_{ij} is the change in specific thermal energy of node i due to its thermal interaction with node j , then the pairwise contribution to the total change in thermal energy is

$$\Delta E_{ij}^{thermal} = m_i \Delta u_{ij} + m_j \Delta u_{ji} \quad (3.87)$$

$$= -m_i \mathbf{v}_i^{n+1/2} \cdot \mathbf{a}_i^n \Delta t - m_j \mathbf{v}_j^{n+1/2} \cdot \mathbf{a}_{ji}^n \Delta t \quad (3.88)$$

$$= m_i \left(\mathbf{v}_j^{n+1/2} - \mathbf{v}_i^{n+1/2} \right) \cdot \mathbf{a}_{ij}^n \Delta t \quad (3.89)$$

where we have used the fact that $m_i \mathbf{a}_{ij} = -m_j \mathbf{a}_{ji}$. Now that we have related the pairwise thermal energy change to the hydrodynamic contributions to nodes from their neighbors, we are free to distribute heat between any two thermally interacting nodes such that the thermal contribution to node i due to j is

$$\Delta u_{ij} = f_{ij} \frac{\Delta E_{ij}^{thermal}}{m_i} = f_{ij} \left(\mathbf{v}_j^{n+1/2} - \mathbf{v}_i^{n+1/2} \right) \cdot \mathbf{a}_{ij}^n \Delta t \quad (3.90)$$

where f_{ij} is a function that distributes energy between nodes i and j such that $f_{ij} + f_{ji} = 1$. The job of this function is to divvy up the heat in a physical way that directs more heat to the cooler of the two nodes. Of particular interest are a simple smooth variation-diminishing form

$$f_{ij} = \frac{1}{2} \left(1 + \frac{u_{ji} \operatorname{sgn}(\Delta E_{ij}^{thermal})}{|u_{ji}| + (1 + |u_{ji}|)^{-1}} \right), \quad u_{ji} \equiv u_j - u_i \quad (3.91)$$

and a strict monotonic form

$$f_{ij}^{mono} = \begin{cases} \max(0, \operatorname{sgn}(B)) : & |B| \leq 1 \\ \frac{m_i}{\Delta E_{ij}^{thermal}} \left(\frac{\Delta E_{ij}^{thermal} + m_i u_i^n + m_j u_j^n}{m_i + m_j} - u_i^n \right) : & |B| > 1 \end{cases} \quad (3.92)$$

where

$$A = \frac{\Delta E_{ij}^{thermal}}{u_{ji}}, B = \begin{cases} A/m_i : & A \geq 0 \\ A/m_j : & A < 0 \end{cases} \quad (3.93)$$

Above, we have used the sign function $\operatorname{sgn}(x)$ that returns 1 for positive arguments, -1 for negative arguments, and 0 otherwise.

While the monotonic form is the most robust, preventing spurious cooling, the most accurate form is a hybrid between this and the “all-or-nothing” form. We use this hybrid form when we employ this algorithm,

unless we are forced to use the monotonic form because of spurious cooling (e.g. in cases where u_i drops below zero for some reason). In any case, given the pairwise nodal thermal energy contributions, the thermal energy gained by node i is $\Delta u_i = \sum_j \Delta u_{ij}$.

We have extended this technique for use in magnetohydrodynamic flows, though it is no longer able to exactly conserve energy once the third energy pool— that of the magnetic field—introduces an ambiguity into the transfer process. Regardless, the results are an improvement over the unsymmetric energy equation in magnetized test problems. We discuss this extension in the next chapter.

The greatest constraint of our implementation of the compatible energy algorithm is that it relies on the assumption that the node connectivity remains the same over an entire time step, since it records the accelerations between all of the interacting node pairs. This means that the algorithm cannot be used when neighbor searches are performed more than once per step, as is necessary with multi-stage integrators under certain boundary conditions. However, since we use the single-stage “Leapfrog” integrator, we can use this compatible energy update for all problems of interest.

3.5 Multidimensional SPH

The equations we have been using to describe SPH already work in multiple dimensions. Thus this technique is easier to generalize than many other methods. There are, however, a few practical concerns we must face as we run more realistic simulations.

We confine this section to concerns about multidimensional hydrodynamic flows. In the next chapter, we will see that presence of magnetic fields, poses other formidable multidimensional problems.

3.5.1 Artificial Viscosity

We have already mentioned the difficulty that current SPH artificial viscosity models face in distinguishing between shear and compressive flows. Since the idea of “shear” doesn’t make sense in one dimension, this is an inherently multidimensional issue.

3.5.2 Computational Expense

Because the aforementioned stencil of SPH is so dynamic, the method suffers somewhat more than others in 2D and 3D, where the number of neighbors of a node can grow significantly, depending on n_h . Contrast the 30-50 points of a 2D SPH stencil with the 4-20 points of a 2D finite difference method. In 3D, the difference is even more pronounced, with an SPH stencil having potentially hundreds of neighbors while that of finite difference method is still in the tens. The fact that a neighbor search must be performed one or more times per step exacerbates the issue.

This expense can be mitigated somewhat by the Lagrangian frame in which SPH operates. For problems in which the mass is relatively localized, a great deal of time can be saved because the quiescent regions are not represented numerically at all.

3.6 Boundary Conditions

Because we chose to neglect the surface term in (3.10) and its multi-dimensional counterparts, we must take special care to enforce the boundary conditions. If we do nothing, the SPH approximation will underestimate quantities and overestimate gradients because there are no points to be sampled outside the boundary. The most common practice in SPH is to replace the boundary with a distribution of “ghost” nodes with values that attempt to recreate the physics at the boundary. This justifies the aforementioned neglect of the surface term, since there is no longer a surface visible to any points for which we are solving (3.30). This simple idea works well for several boundary conditions of interest and is the approach we adopt in this thesis. Examples of the two most common boundary configurations are shown in Figure 3.4.

Typically, we set up a set of ghost nodes that will act as the boundary, and the properties of the nodes (including their positions, velocities, and the like) have to be updated when they are used in the fluid equations. In practice, we can get good results by updating the ghost nodes at the beginning of a simulation step, but this is not always the case.

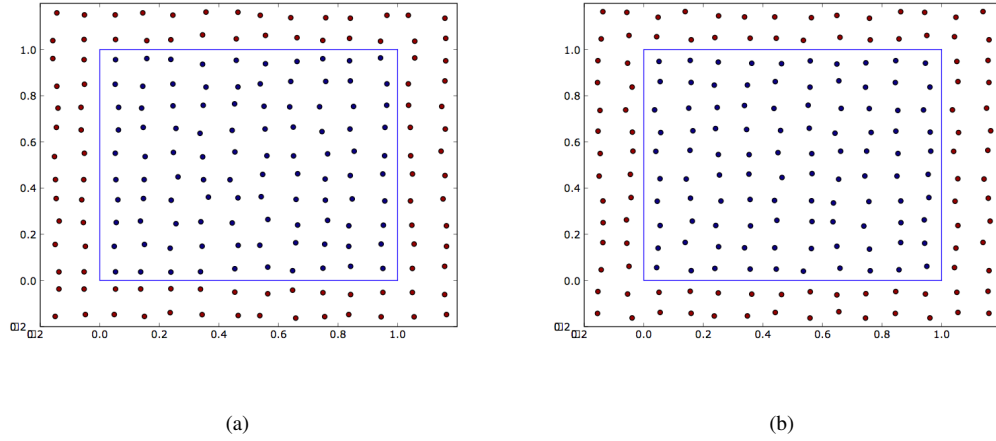


Figure 3.4: Ghost nodes for a set of interior nodes for (a) rigid and (b) periodic boundaries.

3.6.1 Reflecting and rigid boundaries

One of the simplest and most useful boundary conditions one can fathom for hydrodynamic calculations is an impermeable rigid wall that confines a fluid, reflecting any incoming nodes off its surface. This effect can be achieved with admirable precision by creating an “image” node for each interior node whose position and velocity are the reflection of that of the original. The ghost nodes for a set of interior nodes on a domain $[0, 0] \otimes [1, 1]$ with rigid boundaries is shown in Figure 3.4(a).

If an interior node penetrates the surface represented by these image nodes, its position and velocity must be reflected back into the interior. The other physical quantities of the image node may be the same as their interior counterparts, or they may be reflected, depending upon the specific requirements for the boundary.

3.6.2 Periodic boundaries

When we wish to represent a region without boundaries, or a region on the interior of some much larger region, or the transmission of a signal over a very long distance, we can define pairs of periodic boundaries that are physically identified with one another. Consider such a pair consisting of boundaries A

and B . The interior nodes near A are duplicated as ghost nodes behind B , and so forth. Nodes passing into A are emitted by B with the same velocity and other nodal quantities; the node appears displaced from B by an amount equal to which it penetrated A . Figure 3.4(b) shows the ghost nodes created with two pairs of boundaries: one north-south boundary and one east-west.

Periodic boundary conditions work well as long as the nodes are relatively ordered near the boundaries. Small asymmetries and defects in a nodal lattice can often cause significant spurious signals. For this reason, ghost nodes representing periodic boundaries should be updated more often so that these defects can be tempered. In other words, time integration schemes with multiple stages should update the values of periodic ghost nodes at each stage, not merely at the beginning of a step.

3.6.3 Inflow and outflow boundaries

Some of our test problems involve fluids that flow in from or out of boundaries, and so the number of nodes representing these fluids actually changes over the course of the simulation. Such *inflow* and *outflow* boundaries are impractical to represent in the most general sense, since we cannot know what lies outside of the computational domain in any more detail than we explicitly state. Our treatment of these boundaries assumes that

- the density of the incoming/outgoing fluid is uniform in the vicinity of the boundary
- the flow velocity is constant near the boundary
- only one type of material flows through each boundary

These constraints allow us to specify a uniform nodal spacing of nodes passing into and out of the computational domain.

Near inflow boundaries, the initial spacing of the interior nodes Δx is determined, assuming that the nodes are uniformly spaced and that the nearest nodes are placed at a distance $\Delta x/2$ from the boundary. Then ghost nodes are planted behind the inflow boundary in the same manner and with a velocity matching that

of the interior nodes. When the ghost nodes are updated, those within $\Delta x/2$ of the boundary become interior nodes, and the rest are recreated from the new set of interior nodes based upon their initial spacing.

Near outflow boundaries, the initial spacing of the interior nodes is determined in the same manner as for inflow boundaries. When an interior node flows through the boundary, it is removed and replaced with an identical ghost node. The ghost nodes are then updated based upon the positions of the nearest interior nodes and their initial spacing.

3.6.4 Dirichlet boundaries

Sometimes we wish to prescribe the properties of the fluid at the boundary. This type of case is typically called a *Dirichlet* boundary condition by those in the finite difference community and an *essential* boundary condition by the practitioners of finite elements. We have found no mnemonic name for this boundary condition in the literature.

This boundary condition is more often applied to electromagnetic fields than hydrodynamic quantities. In this case it is often safe to assume that the flow velocity is small, and ghost nodes may be placed outside the boundary at desired locations without accounting for significant motion. Then each node's properties are set according to the solution prescribed at its location.

This is a simple matter if the solution is prescribed in the form of an analytic function. If the solution is known only at the boundary then some mathematical handiwork (a Taylor expansion, perhaps) may be done to extrapolate it to the positions of the ghost nodes. If the flow velocity is significant near one of these boundaries, care must be taken in the placement of the ghost nodes.

3.6.5 Neumann boundaries

Also named after one of its inventors, a Neumann (or *natural* in the Finite Element community) boundary condition specifies the derivative of a quantity at the boundary instead of the value of the quantity itself. A scalar quantity may have its normal derivative specified at the boundary, whereas a 3D vector field may be constrained in terms of its curl.

This poses a trickier problem for SPH than it does for grid-based methods, since the topology of the nodes (and hence the definition of spatial derivatives) is much more complicated. For example, the SPH gradient at a point \mathbf{x} is

$$\langle \nabla f(\mathbf{x}) \rangle = \sum_j \frac{m_j}{\rho_j} f_j \nabla W(\mathbf{x} - \mathbf{x}_j, h(\mathbf{x})) \quad (3.94)$$

If we wish to specify that $\mathbf{n} \cdot \nabla f = g$ at the boundary, we must solve a linear system expressing

$$\mathbf{n} \cdot \langle \nabla f(\mathbf{x}_i) \rangle = g(\mathbf{x}_i) \quad (3.95)$$

for a set of nodes $\{\mathbf{x}_i\}$ on the boundary. The unknowns in the system are the values of f on boundary nodes that have been reflected across the surface.

In practice, it is difficult to accurately enforce boundary conditions of this type. Fortunately, this particular boundary condition gives us some insight into handling free boundaries, which we discuss next.

3.6.6 Free boundaries

When a portion of a fluid is out of contact with a rigid boundary, we say that portion constitutes a *free* boundary. In this case, the “outermost” nodes in a distribution represent the boundary of a fluid that moves freely in space depending upon the motion of the fluid. A reliable way for identifying these *free boundary nodes* is to compute the SPH *first moment*

$$\mathbf{I}_1 = \int_{\Omega_i} \mathbf{x}' W(\mathbf{x} - \mathbf{x}'; h) d\Omega \approx \sum_j \mathbf{x}_j \frac{m_j}{\rho_j} W_{ij}(h_i) \quad (3.96)$$

and mark those for which $|\mathbf{I}_1|$ exceeds some threshold value. The idea here is that $\mathbf{I}_1 = 0$ for all nodes on the interior, and so nodes with an appreciable first moment are near or on the boundary.

Once the free boundary nodes are identified, relevant boundary conditions can be applied by solving a linear system for field values on these nodes similar to the method described for that of Neumann boundaries, above. The difference is that, in this case, the unknowns are the free boundary nodes themselves

and not boundary nodes that have been reflected across a mathematical surface. In the next chapter, we will discuss how this free boundary condition helps us represent the interface between a plasma and vacuum.

3.6.7 Cylindrically symmetric systems

The use of cylindrical coordinates is prevalent in the study of magnetized plasmas because the most interesting tractable configurations of confinement have cylindrical symmetry. There have been a few attempts to model cylindrically-symmetric systems in two-dimensions using SPH (see, for example, [113]), but the behavior near the z axis ($r = 0$) remains an issue. Moreover, implementations of these techniques are not widely available and we cannot use them with the same degree of confidence as we can the traditional Cartesian formulation. Are we then forced to run three-dimensional simulations for this important case that is inherently two-dimensional?

The answer depends on whether the region near the axis of symmetry is modeled. There is a relatively easy trick we can use to treat a cylindrically-symmetric system with a number of nodes that scales as N^2 instead of N^3 , provided that we stay away from the z axis. By assuming that $\frac{\partial}{\partial \phi} = 0$ for all quantities of interest, we can represent the interior of such a system as a cross section at $\phi = 0$.

We wish to create ghost nodes identical to these interior nodes about the azimuth so that the interior nodes will see their contributions. We determine an asymuthal spacing $\Delta\phi$ based upon the radial spacing Δr of the interior nodes. If we declare that $r_i \Delta\phi \approx \Delta r$, then $\Delta\phi = \Delta r / r_i$ for interior nodes with $r = r_i$. Ghost nodes can then be created so that there are n_h asymuthal nodes to either side of the interior node at $\phi = 0$. When the ghost nodes are updated, their properties become identical to those of their corresponding interior node.

The problems with this approach near the axis are obvious. As $r \rightarrow 0$, the spacing $\Delta\phi$ becomes large, and interior nodes near the z axis can see an entire ring of ghost nodes around the axis. This occurs because we have imposed a rigid structure on the placement of the azimuthal ghost nodes. In a fully three-dimensional simulation, we would place the nodes in a more stable configuration.

A more robust approach is to divide the azimuth into a number of wedges (2, 4, or 6, perhaps) and place the nodes so that the distribution is symmetric when nodes are reflected across the dividing planes.

Then only one wedge need be simulated in order to represent the entire problem, and ghost nodes can be used to represent nodes visible to the nodes in this wedge. In this approach, it is necessary to duplicate nodes along each wedge's boundary, since all these nodes become visible near the z axis.

If nodes are placed on a uniform cubic lattice, we can easily divide the azimuth into 2 or 4 wedges without affecting the node distribution. For a distribution that uses the close-packed hexagonal lattice, 6 divisions would preserve axisymmetric symmetry.

The challenge of this technique is to ensure that any field values are symmetric under rotations of $2\pi/n$, where n is the number of wedges. Thus, scalar fields remain invariant under this rotation, while vector fields are rotated so that their expression in cylindrical coordinates remains invariant.

Both of these techniques are sensitive to the placement of nodes near the wedge boundaries, so it is important to impose the boundary conditions on both the ghost node positions and their field values at each stage of a multi-step time integration scheme, as is recommended with the periodic boundary condition. This is hardly surprising, since this boundary condition enforces periodic symmetry of order n about the axis.

3.7 Curved Surfaces

Most of the aforementioned boundary conditions have been explored by various authors in the SPH literature using planar surfaces. This is because it is easy to place “ghost” nodes such that desired properties are attained at a plane. It is less obvious to decide how to place ghost nodes to satisfy boundary conditions on a curved surface. Yet many interesting problems in the laboratory require that we be able to represent cylindrical, spherical, and even conical surfaces. A few authors have explored boundary conditions on curved surfaces in SPH, but these techniques are far from general in comparison to flexibility of the planar surface treatment.

3.7.1 Rigid Body Treatment

One general approach to modeling curved surfaces is to model them with nodes representing solid materials instead of representing them with surfaces at the boundary. This can be accomplished “naively” by adding material strength to the fluid equations, using a stress tensor σ in the momentum equation instead of a scalar pressure:

$$\frac{D\mathbf{v}}{Dt} = \frac{\nabla \cdot \sigma}{\rho} \quad (3.97)$$

as in [114]. Thus the fluid equations become the equations of *elastodynamics*. Indeed, SPH has been used successfully for several such problems [115] and so this should work. But if the body is truly rigid (especially when compared to diffuse gas or plasma), this method is a lot of extra work. Moreover, the sound speed c_s in a rigid body is very high: the more rigid the body, the faster that sound waves will travel through its crystalline structure only to bounce back with a force equal and opposite to that which created them, since $c_s = \sqrt{\frac{\partial p}{\partial \rho}}$. And since the CFL condition (3.64) for the time step size varies inversely with this sound speed, simulations will proceed with small steps. Perhaps the most frustrating aspect of this approach is that any deformation in the body resulting from solving the elastic equations is an error for a rigid body!

As a shortcut, Monaghan [116] and others ([117], [118]) have used SPH to represent non-deformable solid bodies with sets of boundary nodes. In these approaches, a solid body contains no nodes inside its surface; if allowed to move, it translates and rotates without deformation, so these nodes would remain at a fixed displacement with respect to the center of mass of the object. These are truly *rigid* bodies in which the deformation gradient $\nabla \mathbf{v}$ is exactly zero. Since rigid bodies may have internal magnetic fields, we have extended this method to include interior nodes. Such an extension is also necessary if one wishes to solve heat transfer problems involving both rigid bodies and fluids.

In general, a rigid body accelerates linearly in response to the net force applied to it and rotationally in response to the net torque. However, we are only interested in the motion of the fluid colliding with a stationary body. It is difficult to use the conservation of momentum to represent such collisions in terms of

opposing forces: these forces are exerted over the time interval of the collision, so representing instantaneous collisions is problematic.

Monaghan solves this problem by expressing the force of a rigid body upon a deformable material as a Lennard-Jones force [116] reminiscent of those used in molecular dynamics. The force exerted on a node i within a deformable body by a node j within a rigid body takes the form

$$\mathbf{F}_{ij} = D \left[\left(\frac{r_0}{|\mathbf{x}_{ij}|} \right)^{p_1} - \left(\frac{r_0}{|\mathbf{x}_{ij}|} \right)^{p_2} \right] \frac{\mathbf{x}_{ij}}{|\mathbf{x}_{ij}|^2} \quad (3.98)$$

where D , p_1 , and p_2 are fit parameters and r_0 is a measure of the spacing between nodes in the rigid body. Typically, $p_1 = 4$ and $p_2 = 2$, and D can be set according to material properties or simply set constant.

This method has been used to study water waves created by a falling solid [119], swimming linked bodies [118], and violent ocean waves [117]. These problems have flows at relatively low velocities, at which the Lennard-Jones method yields good results. Unfortunately, we find that this technique produces explosive instabilities when fluids at higher velocities impact the solids. Other authors have reported the same instabilities [120]. Our interests involve plasmas at high velocities, so we cannot rely on it.

3.7.2 The Single and Multiple Tangent Boundary Methods

Some authors have developed a technique closer to the traditional plane-reflection method in modeling incompressible flows about curved surfaces. Morris et al [121] describe a method in which a normal distance is computed for each interior node, and that distance is used to construct a tangent plane across which nodes are reflected. Yildiz et al [120] expand upon this method by creating several tangent planes by choosing several control points on the curved boundary and repeating Morris's procedure for each control point with respect to the interior nodes. Yildez dubbed this the "multiple boundary tangent method" to Morris's "single boundary tangent method."

These tangent-plane-based methods are relatively simple and effective and have been used to model flow past obstacles at low Reynold's numbers. Unfortunately, in the case of the multiple boundary tangent method, one must predefine control points on curved boundaries in order to construct the tangent planes used to create

ghost nodes. Also, one must take care to divide the mass evenly between ghost nodes created from nodes reflected across multiple tangent planes.

3.7.3 Simple Reflection and Scaling Method

A simpler (but less accurate) procedure to model certain curved boundaries is to express the normal vector of the boundary analytically and to compute the normal distance of each interior node near the boundary accordingly. Ghost nodes are then constructed by reflecting interior nodes across this surface as usual. This is a less general method than the tangent-plane-based ones, since one must know the equation for \mathbf{n} everywhere on the boundary, but we currently are most interested in simple curved surfaces like cylinders, cones, and spheres.

Reflecting interior nodes across the surface is not sufficient to provide accurate boundary conditions, because there will invariably be too many or too few ghost nodes. If the “interior” of a domain lies within the surface, there will be a lower density of ghost nodes outside, and if the “interior” is outside the surface, there will be a higher density inside. To counteract this imbalance, we compute the nodal masses and SPH smoothing scales of the ghost nodes by scaling those of the interior nodes.

Imagine a single layer of ghost nodes of uniform mass m_g on the outside of a cylinder Γ of length L and radius R . The ghost nodes are all placed at a radius r_g from the center of Γ and represent a fictitious shell of material. The shell has a volume approximately equal to $\pi L(r_g^2 - R^2)$. Likewise, suppose there is a shell of interior nodes with uniform mass m_i within Γ from which the ghost nodes were created. If these nodes are placed at a radius r_i from the Γ ’s center, this shell has an approximate volume of $\pi L(R^2 - r_i^2)$. If the mass density is to be equal across the surface, the ratio of the masses of these shells must be equal the ratio of their volumes. The ratio of the shell masses is the ratio of the interior and ghost nodal masses:

$$\frac{m_i}{m_g} = \frac{R^2 - r_i^2}{r_g^2 - R^2} \quad (3.99)$$

Likewise, since the volume of a v -dimensional node is roughly equivalent to h^v , we can relate the SPH smoothing scales h_i and h_g by

$$\frac{h_i}{h_g} = \frac{R^2 - r_i^2}{r_g^2 - R^2} \quad (3.100)$$

3.8 Test Problems

Here we describe problems are used to verify that the techniques we have discussed can treat the types of problems in which we are interested. Where possible, we compare our results with known solutions and measure the rate of convergence as the resolution is increased. Specifically, we measure the L_1 , L_2 , and L_∞ norms of the error in the computed solution. These norms are defined as

$$L_p(q) = \left(\frac{\sum_{i=1}^N q_i - q_i^e}{N} \right)^{1/N} \quad (3.101)$$

$$L_\infty(q) = \max_{1 \leq i \leq N} |q_i - q_i^e|$$

where q is the quantity under consideration, q_i is its computed solution at the point \mathbf{x}_i , and q_i^e is the corresponding exact solution. These norms are computed for various resolutions, and the convergence rate A_p of the L_p norm is obtained by computing the successive ratios

$$X = \log\left(\frac{N}{N_0}\right) \quad , \quad Y_p = \log\left(\frac{L_p}{L_{p0}}\right) \quad (3.102)$$

and fitting them to the line $Y_p = -A_p X + B_p$. Above, N_0 is the lowest resolution and L_{p0} is the associated L_p norm.

The use of convergence studies is less common than it should be in the SPH community. In fact, most of the papers we cite from this community do not contain rigorous convergence studies. Those that do probe different resolutions tend to describe loosely the trends that they see in certain quantities. Understandably, this has fostered much skepticism in the mesh-based community regarding the reliability of SPH([100], [122]). The lack of analytic solutions is a concern for studies of convergence, but Stone points out that this methodology can also be applied to experiment [72]. Therefore, one goal of this thesis is to

better quantify the accuracy of SPH and its refinements for magnetohydrodynamic problems. Our convergence studies will usually be shown with plots of error norms, annotated by the rate of convergence, r , of each norm, along with any relevant error bounds.

In our simulations, we use the B-spline kernel defined by (3.3). This kernel has an extent $\kappa = 2$ and is a very popular choice among SPH authors. Unless otherwise specified, we use the Leapfrog integrator described in Section 3.3. We also typically use the XSPH velocity smoothing algorithm with a factor of 1, since it tends to prevent unwanted node interpenetration. For artificial viscosity, we use the Monaghan-Gingold model with linear coefficient $\alpha = 1$ and quadratic coefficient $\beta = 3/4$. Unless otherwise specified we use a CFL safety factor of 0.5.

3.8.1 Hydrostatic equilibrium

The simplest test of any numerical time-dependent system is whether it can maintain equilibrium. In the absence of gravity, a fluid is in hydrostatic equilibrium if $\nabla p = 0$ everywhere. For all of the configurations we describe, we place nodes in such a way that the fluid is uniformly dense everywhere (or at least approximately so) and use an ideal gas equation of state so that the pressure is also constant. We use this test mainly to ensure that our boundary conditions behave properly in the simplest case.

The most obvious version of this test is to lay down nodes on a periodic uniform lattice in 1, 2, and 3 dimensions. Both rigid and periodic boundary conditions are used. Equilibrium is tested on domains with nodal spacings Δx of 0.005, 0.0025, 0.00125, and 0.000675 over the time interval $[0, 1]$. The number of nodes per dimension within the repeated domain is $4 \times \kappa \times n_h$, where κ is the extent of the smoothing kernel W as measured in smoothing scales.

We begin by studying the effect of n_h on this equilibrium using the unsymmetric form of the energy equation and the compatible energy formulation. Most SPH authors use values of n_h between 1 and 1.2. Price tends to choose values between 1.2 and 1.4. Table 3.3 shows the L_p error norms of the density for values of n_h between 1 and 2 in 1D simulations using the unsymmetric SPH energy equation and the compatible energy algorithm. There we have listed only the L_1 norms because the simulations that do converge in 1D produce

identical L_1 , L_2 , and L_∞ norms that do not vary with resolution.

n_h	$L_1(\rho)$
1	-
1.2	0.0018
1.4	0.0044
1.6	0.0026
1.8	0.00035
2	9×10^{-7}

Table 3.3: $L_1(\rho)$ in 1D hydrostatic equilibrium for various n_h

Simulations in which $n_h = 1$ are not convergent. In 1D, the system dismantles itself completely in the unsymmetric form, and produces negative thermal energies in the compatible energy form (the unstable configuration produces large velocities that rob the nodes of their heat). If the simulation is run out to $t = 0.1$, the simulation completes, but we measure largely negative convergence rates in ρ and \mathbf{v} .

At $n_h = 1.2$ and above, the 1D results are better. There is measurable distortion in the density profile, fluctuating through the n_h parameter space but becoming insignificant at $n_h = 2$. In any case, the nodal acceleration is essentially zero (around 10^{-16}).

The results in 2D and 3D are more interesting. At $n_h = 1$ and $n_h = 1.2$ the system is completely unstable and the simulations do not finish because of nontrivial fluid motion. In fact, $n_h = 1.4$ does not even finish for 3D simulations. Simulations with higher values of n_h do finish successfully, but with mixed degrees of success. The “convergence” of the error norms for the velocity is very volatile, but the simulations typically produce small nodal velocities (on the order of 10^{-6} and below) when n_h is sufficient for the simulation to reach $t = 1$.

In any case, there is significant motion in the nodes in 2D and 3D. Figure 3.5 shows an example of this motion in 2D. This “thermal” motion is discussed at length by Martin et al [82], who notes that a cubic lattice is not a numerically stable configuration for SPH. He recommends letting nodes “settle” for several steps before beginning a simulation. The stability of SPH to the initial node configuration is evidently a strong function of n_h , and the lower values of this parameter most commonly in use simply do not yield convergent

results in 2D and 3D.




Figure 3.5: The positions of nodes initially in a cubic lattice after thermal motion with $n_h = 1.8$. The interior nodes seem to have settled into a close-packed lattice.

3.8.2 An acoustic wave

The simplest solution to our fluid equations is a linear compressional wave propagating across a fluid. This problem is treated in every text on compressible fluid dynamics and is a good indication of the accuracy of a numerical technique in representing sound waves. Though simple conceptually, acoustic waves are not trivially modeled—naïve approaches can lead to waves with the wrong propagation speed, or with significant distortion.

We demonstrate the propagation of an acoustic wave in a perfect gas ($\gamma = 5/3$) along a periodic domain $[0, 1]$ with initial conditions and solutions of the following form, which can be found in countless texts on hydrodynamics [64]:

$$\rho(x, t) = \rho_0 + \rho_1 \sin[k(x - c_s t)] \quad (3.103)$$

$$v(x, t) = \frac{\rho_1}{c_s \rho_0} \sin[k(x - c_s t)] \quad (3.104)$$

$$p(x, t) = (\gamma - 1)\rho_0 u_0 + c_s^2 \rho_1 \sin[k(x - c_s t)] \quad (3.105)$$

$$u(x, t) = \frac{p(x, t)}{(\gamma - 1)\rho(x, t)} \quad (3.106)$$

where k is the wave number of the disturbance, c_s is the sound speed of the medium, ρ_1 is a disturbance in the density, which has an “ambient” value of ρ_0 (with $\rho_1 \ll \rho_0$), v_1 is a velocity perturbation, and u_0 is the ambient specific thermal energy. For this test, we set $u_0 = 1/(\gamma(\gamma + 1))$ so that the sound speed c_s is 1. This means that the wave will cross the domain every 1 time unit. We can then gauge the quality of the signal after an integral number of crossings by comparing it to the initial conditions.

There are two obvious ways to initialize the density profile. The first is to vary the masses of the nodes according to the profile and space the nodes evenly; the second is to use nodes of equal mass and space them in such a way that the density profile is achieved through the summation formula (3.7). The first method is simple, but doesn’t correspond to the physical distribution of matter. In fact, the fact that there are inseparable chunks of “heavier” and “lighter” fluid can create aphysical artifacts in the fluid motion. Therefore we use the second method: the nodal mass m_0 is computed by M/N , where M is the total mass given by

$$M = \int \rho(x) dx = \rho_0 \left[x - \frac{\rho_1}{\rho_0} \frac{\cos[kx]}{k} \right]_0^1. \quad (3.107)$$

Then M_i , the cumulative mass from 0 to x_i , is $\sum_{j < i} m_0(j + \frac{1}{2})$. To find the position of the i th node, we solve the nonlinear equation

$$M_i - \rho_0 \left[x + \frac{\rho_1}{\rho_0} \frac{\cos[kx]}{k} \right]_0^x = 0 \quad (3.108)$$

using Newton-Raphson iteration.

In the following tests, we analyze the signal after 1 and 5 crossings. We characterize the solution's quality in terms of the phase error (measured between the nodes of the solution and the initial conditions) and its general stability. The solution is computed domains at resolutions $N = 100$, $N = 200$, $N = 400$, and $N = 800$ for various configurations. We use a wave number of 2π so that a wavelength spans the entire domain, and the amplitude of the density disturbance is $\rho_1/\rho_0 = 5 \times 10^{-3}$. We use a CFL factor of 0.1 instead of the usual 0.5, since this improves the stability of the treatment. Figure 3.6 shows the density profile of the wave after 1 and 5 crossings with the unsymmetric energy equation, and Figure 3.7 shows the convergence of ρ , v_x , and u under these conditions using $n_h = 2$.




Figure 3.6: An acoustic wave after having passed through the domain (a) once ($t = 1$) and (b) five times ($t = 5$), using the unsymmetric energy equation and $n_h = 2$.

We have also run this simulation at lower values of n_h , tabulating the convergence rates in Table 3.4 and Table 3.5. The convergence of the mass density at these lower values is very poor and in some cases nonexistent! That of the velocity is less sensitive but still increases with n_h on the whole.

We emphasize at this point that these runs use artificial viscosity with the usual settings. We observe significant nonlinear steepening from the viscosity, and this steepening severely limits the convergence of the problem. It is remarkable how much n_h seems to affect the convergence of the solution under these constraints. If no artificial viscosity is used, we do not observe such steepening, but the simulations fail to




Figure 3.7: The convergence rates of (a) the mass density and (b) the velocity of an acoustic wave having crossed a domain five times using the unsymmetric energy equation and $n_h = 2$.

n_h	$r\{L_1(\rho)\}$	$r\{L_2(\rho)\}$	$r\{L_\infty(\rho)\}$
1.2	0.06935 ± 0.0078	0.1257 ± 0.0059	0.1944 ± 0.0055
1.4	$-0.00002071 \pm 1.9e-9$	0.03486 ± 0.00043	0.1854 ± 0.00069
1.6	0.008805 ± 0.00014	0.04109 ± 0.00087	0.06059 ± 0.022
1.8	0.5282 ± 0.0089	0.5271 ± 0.0087	0.5213 ± 0.031
2.0	0.7986 ± 0.0012	0.79 ± 0.0059	0.7643 ± 0.015

Table 3.4: Convergence rates for L_p norms of the mass density in the acoustic wave as functions of n_h .

n_h	$r\{L_1(v)\}$	$r\{L_2(v)\}$	$r\{L_\infty(v)\}$
1.2	0.3059 ± 0.019	0.3007 ± 0.02	0.3216 ± 0.019
1.4	0.7111 ± 0.0024	0.6986 ± 0.0015	0.6727 ± 0.0019
1.6	0.136 ± 0.0046	0.1508 ± 0.0085	0.1709 ± 0.0061
1.8	0.5897 ± 0.014	0.5878 ± 0.0059	0.589 ± 0.00085
2.0	0.799 ± 0.001	0.7911 ± 0.0056	0.7668 ± 0.015

Table 3.5: Convergence rates for L_p norms of the velocity in the acoustic wave as functions of n_h .

complete at higher resolutions. Density profiles of these undamped simulations are shown in Figure 3.8. Note the oscillations developing at the peak of the wave at $t = 5$.




Figure 3.8: An acoustic wave *without artificial viscosity* after having passed through the domain (a) once ($t = 1$) and (b) five times ($t = 5$), using the unsymmetric energy equation and $n_h = 2$.

Following Price's lead [84], we take interest in the effects of improved energy conservation on the solution. However, with the error dominated by excessive steepening, neither the ∇h corrections and the compatible energy method improve the convergence of ρ or \mathbf{v} at all. If we temper the artificial viscosity, zeroing the nonlinear coefficient ($\beta = 0$) and setting the linear coefficient $\alpha = 1/4$, we greatly reduce this steepening and see significant improvements in the convergence in the simulation (see Figure 3.9 and Figure 3.10).

With these tempered coefficients, the error in the solution is no longer dominated by steepening, and we might expect the ∇h terms and/or the compatible energy method to make more of a difference in the solution. Unfortunately, this appears not to be the case—simulations with both of these improvements make no detectable difference in the rates of convergence of the mass density or the velocity. This seems to contradict the findings of Price [90] that the wave propagation speed is improved by the ∇h terms. Of course, it is possible that these corrections would have more of an effect at a lower value of n_h , since the uncorrected algorithm performs quite poorly at values typically used in the SPH community.

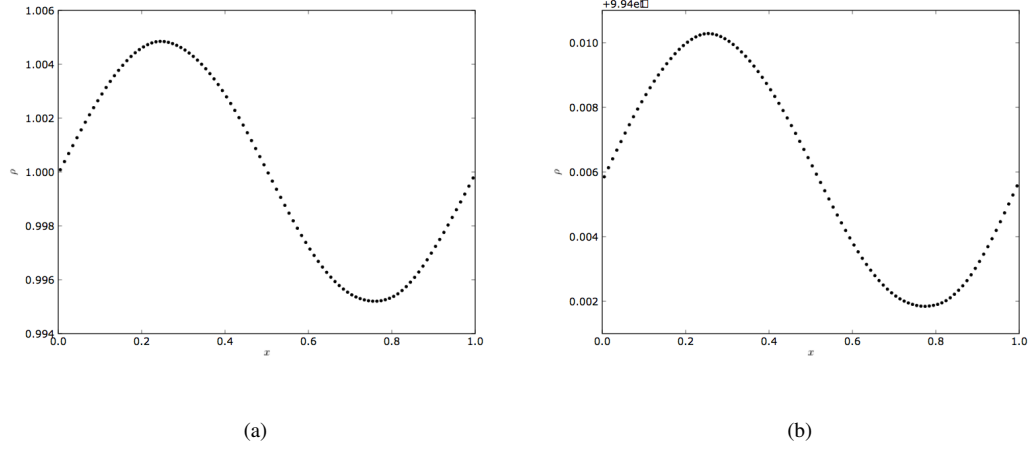


Figure 3.9: An acoustic wave at (a) ($t = 1$) and (b) ($t = 5$) using the unsymmetric energy equation and $n_h = 2$ with tempered artificial viscosity coefficients $\alpha = 1/4$ and $\beta = 0$.

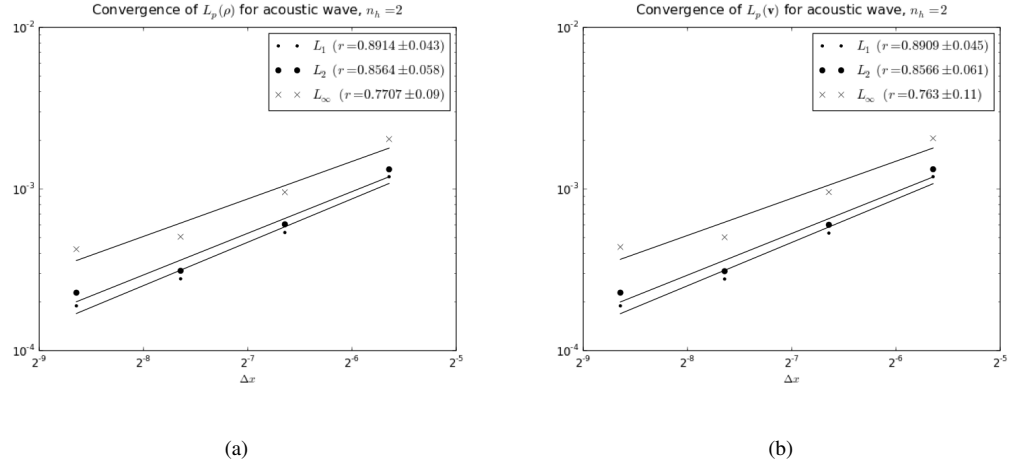


Figure 3.10: The convergence rates of (a) the mass density and (b) the velocity of the acoustic wave at $t = 5$ using the unsymmetric energy equation and $n_h = 2$ with tempered artificial viscosity coefficients $\alpha = 1/4$ and $\beta = 0$.

The results of this test problem are both encouraging and worrying. Clearly, the solution can be made to converge at near first-order, provided one uses a sufficiently robust value of n_h . However, since this problem falls into the linear regime of hydrodynamics, one might expect to achieve something closer to second order, the optimal convergence for SPH. If the problem could be run without artificial viscosity at any resolution, this may very well be the case. However, the advection of such a signal is a problem that is inherently unstable without some form of dissipation. In Eulerian methods, high-order upwinding techniques [123] are used to treat this instability without sacrificing accuracy. Such methods do not currently exist for SPH.

3.8.3 The Sod shock tube

The most famous problem for testing weak hydrodynamic shocks is a 1D shock tube analyzed by Sod [124]. The system consists of a tube spanning $[0, 1]$ containing two perfect gases at rest with $\gamma = 5/3$ separated by a partition at $x = 0$. The gas on the left has initial conditions (ρ_1, p_1) and the one on the right has initial conditions (ρ_2, p_2) , where $\rho_1 > \rho_2$ and $p_1 > p_2$. The partition is lifted and a shock wave propagates to the right, accompanied by a rarefaction to the left. We use the initial conditions (different from those of Sod):

$$\rho_1 = 1 \qquad \rho_2 = 0.125 \qquad (3.109)$$

$$p_1 = 1 \qquad p_2 = 0.1795 \qquad (3.110)$$

First, we compute the solution to this problem with $n_h = 2$ using the unmodified unsymmetric energy equation. The nodes all have equal masses and are spaced so that the summation formula (3.7) is satisfied on either side of the partition at $x = 0$. For the initial conditions (3.110), this means there are four times as many nodes on the left side as there are on the right. We could just as easily vary the masses so that those on the left are four times as heavy as those on the right, but we avoid this for reasons similar to those discussed in the acoustic wave problem. u is computed from the ideal gas relation $p = (\gamma - 1)\rho u$. To ensure that the boundaries of the domain do not contribute to the dynamics, rigid planar boundaries are placed on

either end.

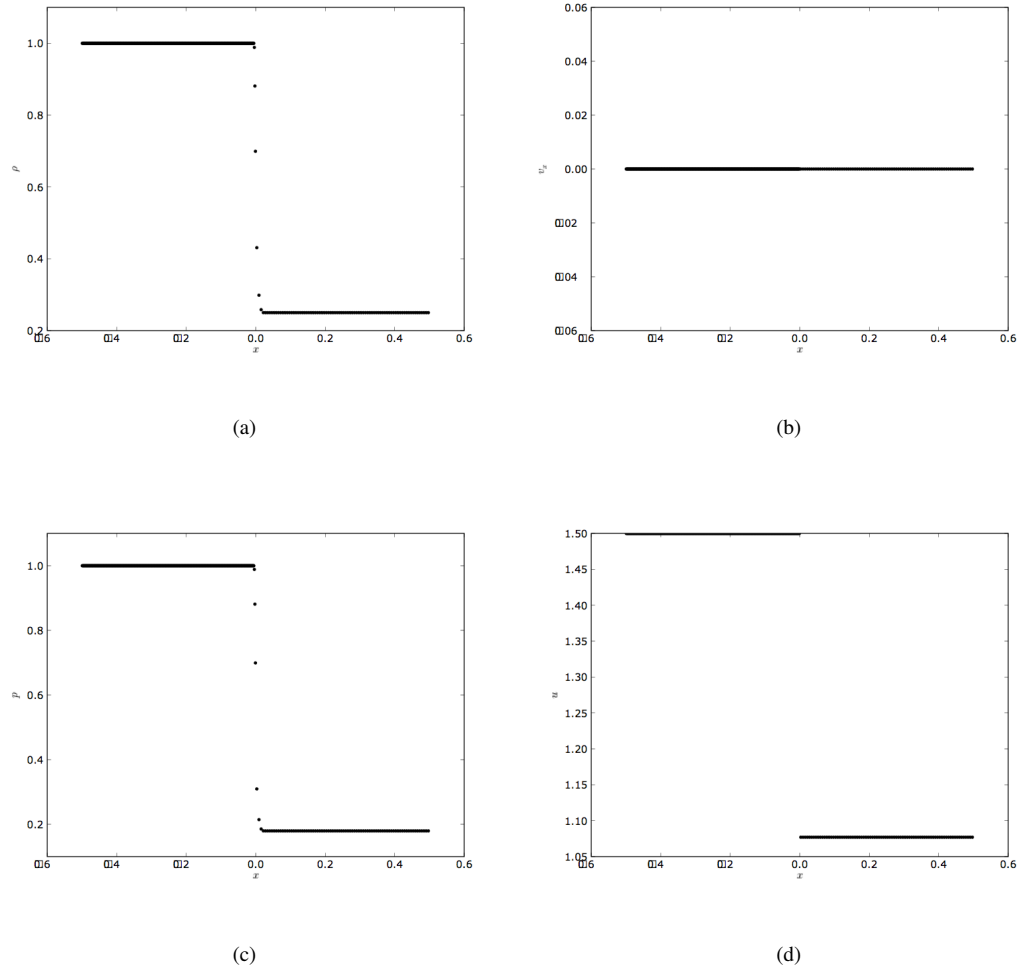




Figure 3.11: The initial conditions for the Sod shock tube problem: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy


The initial and final profiles for ρ , v_x , and u are shown in Figure 3.11 and Figure 3.12, with the SPH nodal values represented by dots and the analytic solutions as solid lines. The shock wave propagates to the right in the velocity and density profiles, while the (continuous) rarefaction from the shock proceeds to the left. In the solution we can see oscillations in the contact discontinuity separating the left and right fluids in the pressure profile. Interestingly, these oscillations do not appear if we use the second-order Runge-Kutta or Predictor-Corrector integrators.




(a)



(b)



(c)



(d)

Figure 3.12: The solution of the Sod shock tube problem at $t = 0.15$ using the unsymmetric energy equation: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy

The convergence of the L_p error norms as defined by (3.101) for simulations at $N = 100$, $N = 200$, $N = 400$, and $N = 800$ are shown in Figure 3.13. Since the problem has discontinuities, we expect first-order convergence at best. We have chosen a high value for n_h : lowering this parameter has effects on these convergence rates that are shown in Table 3.6, Table 3.7, Table 3.8, and Table 3.9. The problem does not run successfully for $n_h = 1$. It is interesting to note that while the convergence rates improve significantly as n_h is increased, $n_h = 1.6$ seems to be a suboptimal setting in the space we have sampled.

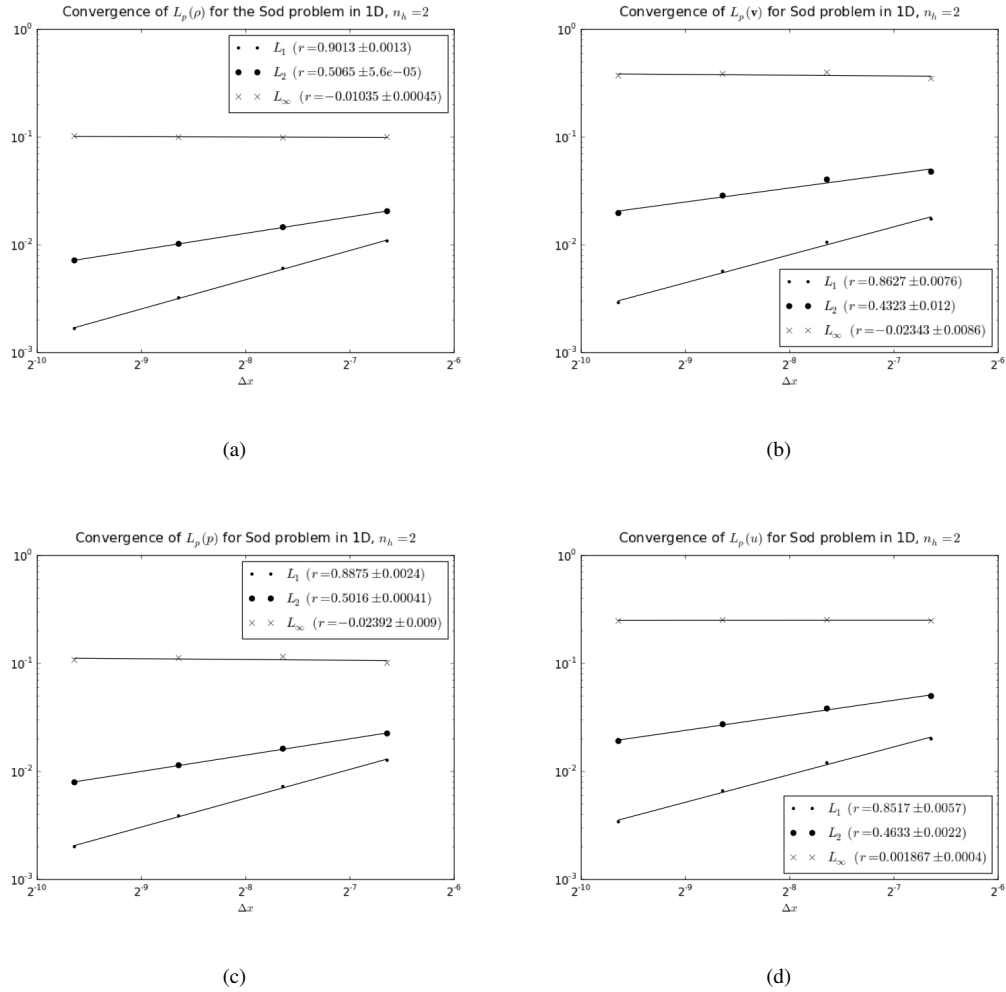


Figure 3.13: Convergence rates for the Sod problem using the unsymmetric energy equation with $n_h = 2$: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy

Since the acoustic wave problem was not affected by the ∇h terms or the compatible energy al-

n_h	$r\{L_1(\rho)\}$	$r\{L_2(\rho)\}$	$r\{L_\infty(\rho)\}$
1.2	0.4209 ± 0.0044	0.3361 ± 0.00016	-0.14 ± 0.0011
1.4	0.4138 ± 0.015	0.4173 ± 0.0015	$-0.0004863 \pm 4.3e-5$
1.6	0.431 ± 0.014	0.4176 ± 0.0017	$0.0113 \pm 1.5e-5$
1.8	0.7871 ± 0.0025	$0.5089 \pm 4.4e-5$	-0.02324 ± 0.0008
2.0	0.9013 ± 0.0013	$0.5065 \pm 5.6e-5$	-0.01035 ± 0.00045

Table 3.6: Convergence rates for L_p norms of the mass density in the Sod problem as functions of n_h .

n_h	$r\{L_1(v_x)\}$	$r\{L_2(v_x)\}$	$r\{L_\infty(v_x)\}$
1.2	0.4593 ± 0.00081	0.2814 ± 0.0073	-0.1514 ± 0.0072
1.4	0.7401 ± 0.0057	0.3633 ± 0.029	-0.1403 ± 0.084
1.6	0.55 ± 0.0042	0.475 ± 0.0083	0.1121 ± 0.023
1.8	0.7781 ± 0.0011	0.4885 ± 0.015	0.03121 ± 0.018
2.0	0.8627 ± 0.0076	0.4323 ± 0.012	-0.02343 ± 0.0086

Table 3.7: Convergence rates for L_p norms of the velocity in the Sod problem as functions of n_h .

n_h	$r\{L_1(p)\}$	$r\{L_2(p)\}$	$r\{L_\infty(p)\}$
1.2	0.4252 ± 0.0035	0.2765 ± 0.00042	-0.2191 ± 0.012
1.4	0.4641 ± 0.012	0.4155 ± 0.00013	-0.1108 ± 0.013
1.6	0.397 ± 0.01	0.3382 ± 0.0036	-0.1532 ± 0.012
1.8	0.7861 ± 0.0013	0.5072 ± 0.00011	-0.03452 ± 0.005
2.0	0.8875 ± 0.0024	0.5016 ± 0.00041	-0.02392 ± 0.009

Table 3.8: Convergence rates for L_p norms of the pressure in the Sod problem as functions of n_h .

n_h	$r\{L_1(u)\}$	$r\{L_2(u)\}$	$r\{L_\infty(u)\}$
1.2	0.5181 ± 0.0024	0.3465 ± 0.0016	-0.1881 ± 0.0096
1.4	0.7837 ± 0.00012	0.4419 ± 0.0011	-0.1222 ± 0.011
1.6	0.5437 ± 0.0066	0.43 ± 0.00051	-0.05698 ± 0.0034
1.8	0.7843 ± 0.00042	0.4858 ± 0.002	0.01283 ± 0.00065
2.0	0.8517 ± 0.0057	0.4633 ± 0.0022	0.001867 ± 0.0004

Table 3.9: Convergence rates for L_p norms of the specific thermal energy in the Sod problem as functions of n_h .

gorithm, we use this problem as a guide to decide how the latter compares to the former in improving the solution. Including the ∇h terms yields the results shown in Figure 3.14. The results using the compatible energy update are shown in Figure 3.15.

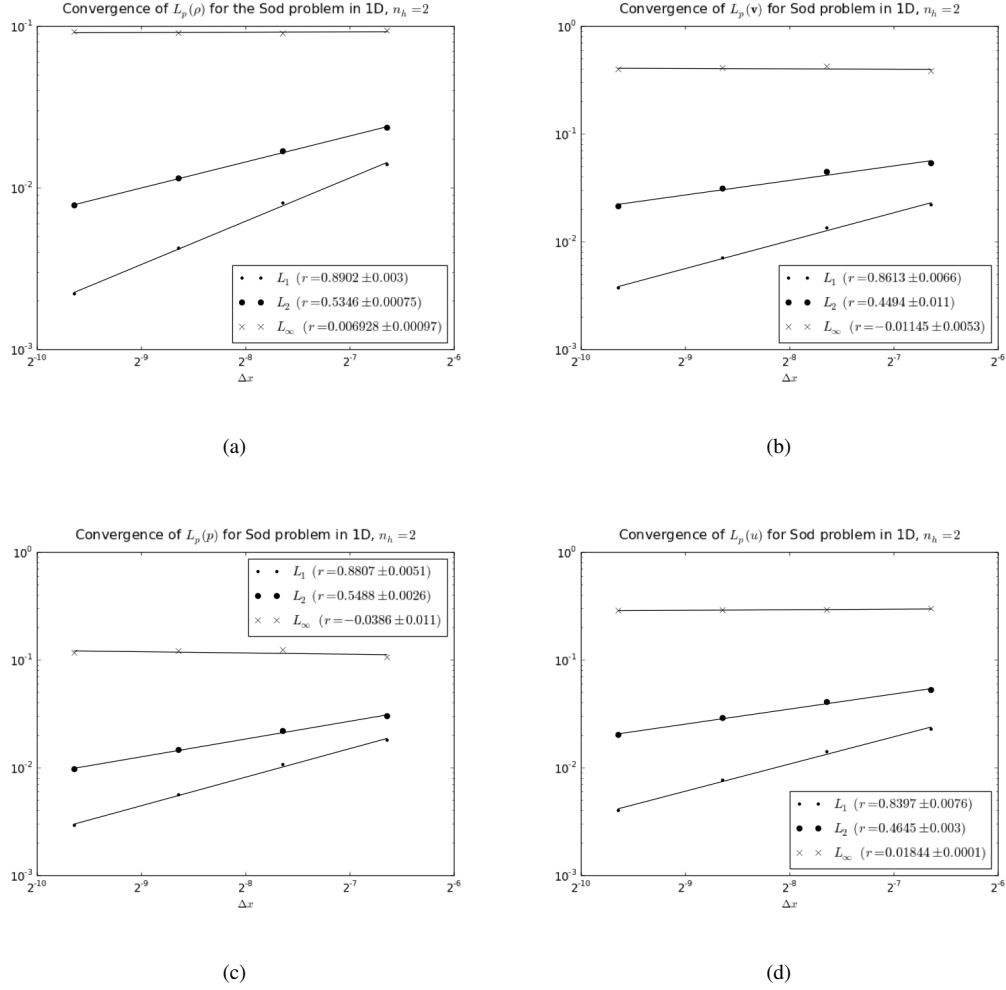


Figure 3.14: Convergence rates for the Sod problem using the unsymmetric energy equation with $n_h = 2$ and ∇h corrections: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy

Both show improved rates of convergence, with the compatible energy method edging out the ∇h improvement in all quantities except (ironically) the specific thermal energy itself. Comparing the specific energy profiles in Figure 3.8.3 demonstrates the difference in the treatment of u between the two refinements. Both curves reproduce the overshoot in the u curve at the contact discontinuity, but the curve in (b) is clearly

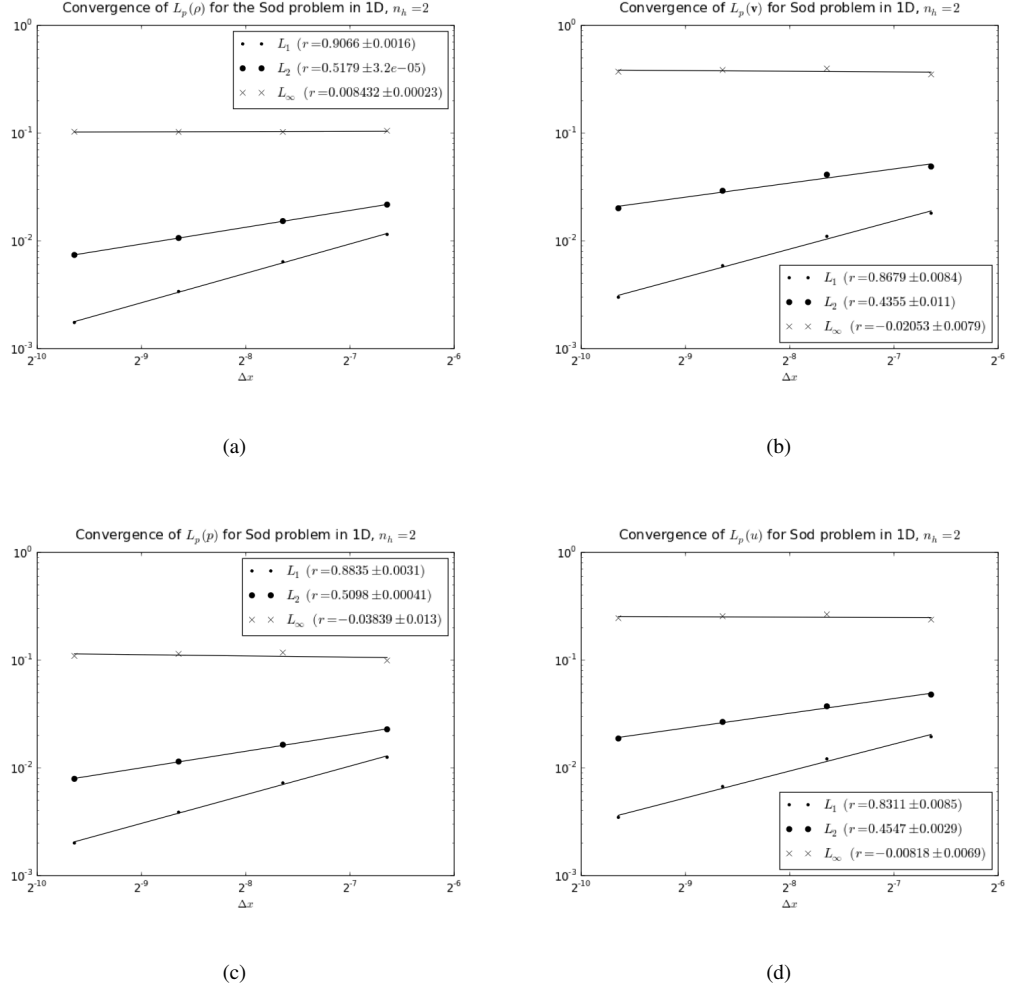


Figure 3.15: Convergence rates for the Sod problem using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy

smoother than that in (a). The smoother curve has more nodes in the neighborhood of the overshoot, worsening the error in that region. Nevertheless, the smoothness in the node distribution is an improvement over that computed by the unsymmetric energy equation with and without the ∇h corrections. In terms of energy conservation, there is no contest between the two methods: the compatible energy method conserves the total energy to machine precision—the other method does not.




Figure 3.16: The specific energy profile for the Sod problem at $t = 0.15$ using (a) ∇h corrections and (b) the compatible energy method.

The ∇h correction factor Ω is very sensitive to the gradient of the smoothing kernel, and it exhibits worrying oscillations even in the presence of relatively weak shocks, as is shown in Figure 3.8.3. When it is used, artifacts in the mass density and the specific thermal energy are a bit more pronounced, perhaps because the accuracy is improved elsewhere, but more likely because Ω changes sign so rapidly at the discontinuities.

The two improvements for the energy equation are certainly not mutually exclusive, and the question naturally arises as to whether it is profitable to combine them. We do this in simulations whose convergence is shown in Figure 3.18. The answer to our question appears to be “no”—the rates of convergence achieved with this combination are worse than the best of either one, and the curve for u is virtually indistinguishable from that computed using the compatible energy method alone.

Now that we are confident that we can solve the Sod problem effectively, we examine some of




Figure 3.17: The ∇h correction factor Ω for the Sod problem at $t = 0.15$.

the other refinements. We have discussed the overshoot in the specific thermal energy that occurs near the contact discontinuity. Let us see whether Price's artificial thermal conduction terms ((3.56)) have any effect. Using the compatible energy method with these terms and a coefficient $\alpha_u = 0.1$, we obtain a profile and a new convergence curve for the specific thermal energy that are shown in Figure 3.8.3. The overshoot in u is smoothed but still present, as is the oscillation in p at the contact discontinuity. The convergence is degraded by the overall smoothing. Cranking α_u higher increases the smoothing at the further expense of accuracy. Price's use of Morris's variable switches does improve the smoothing of u noticeably and may be a viable way for applying artificial thermal diffusion to shocks in the short term. Still, a proper limiter would work much better.

At this point, it is also worth demonstrating the effectiveness of Morris's variable switch for artificial viscosity. Using Price's Riemann-based model ((3.53)) for artificial viscosity with Morris's switch ((3.52)) and the compatible energy method, we obtain the results shown in Figure 3.20, with the convergence rates shown in Figure 3.21.

Clearly, this variable switch allows the emergence of oscillations in the velocity profile, indicating that its criterion for detecting nonlinear regions is too stringent. Given that this is a relatively weak shock problem, this is unsettling. On the other hand, the convergence rates for the quantities are all noticeably

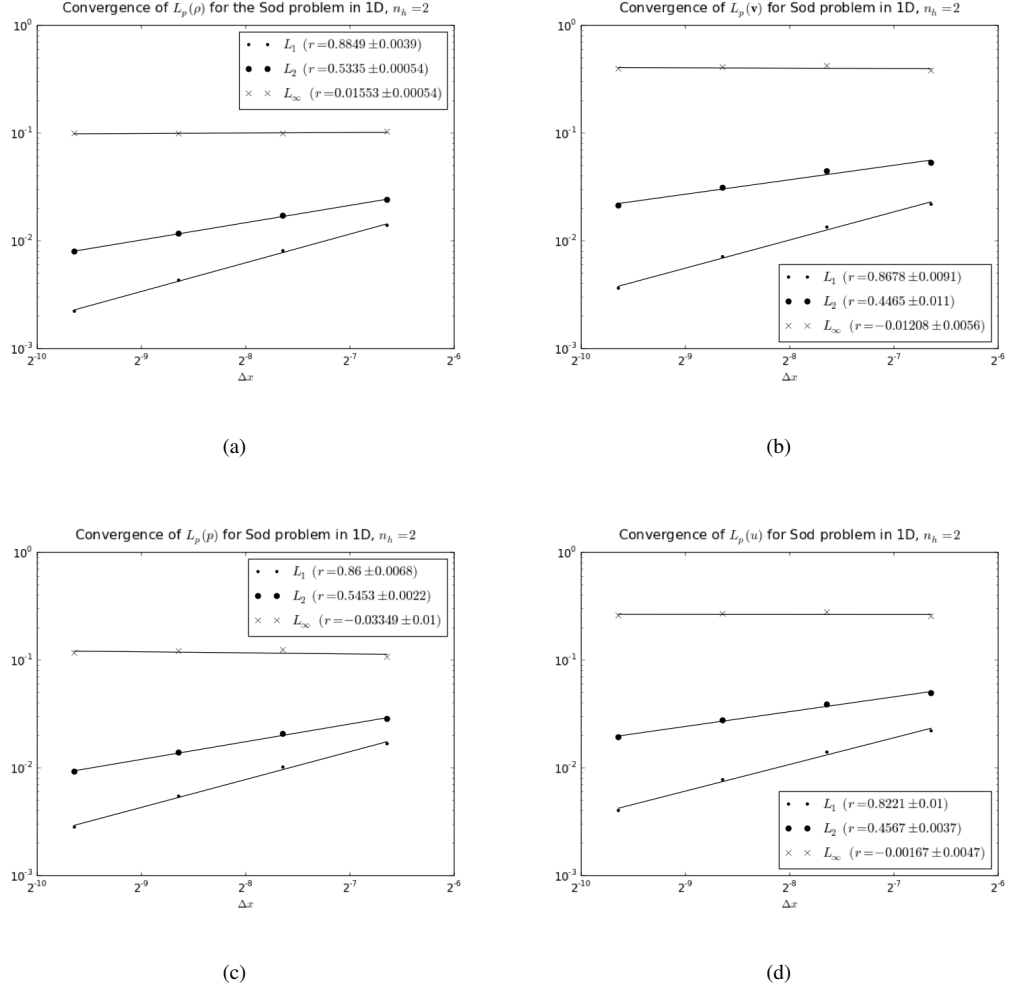


Figure 3.18: Convergence rates for the Sod problem using the compatible energy method with $n_h = 2$ and with ∇h corrections: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy




Figure 3.19: The solution (a) and convergence (b) of the specific thermal energy in the Sod problem at $t = 0.15$ using the compatible energy method and $n_h = 2$ with Price's artificial heat conduction terms.

improved. More than anything else, this indicates that better results can be achieved by tempering the artificial viscosity. However, the oscillations in v_x lead us to discard this variable switch approach in our further simulations. Price also holds the dissipation parameters constant in his studies [84].

Comparison with a Godunov code

In the next chapter, we will be dealing with test problems that do not have known solutions. It is common practice in this case to take as canon the results of a highly-resolved simulation using a trusted numeric method and perform convergence studies with respect to these numeric results. This is far from ideal, but is simply the best option available absent an analytic solution! The favored methods for analyzing shock tube problems are those derived from Riemann solvers, since they are the only methods that manifestly accomodate the shock relations, and (in one dimension, at least) they can justifiably claim to be based off of pertinent analytic solutions.

In preparation for such numeric convergence studies, we now analyze the results of the Sod problem as obtained by Athena [40], a freely-available Godunov grid-based code that solves the ideal MHD equations. To compare our SPH results directly with those of Athena, we interpolate the nodal values of ρ , \mathbf{v} , and u




Figure 3.20: The solution of the Sod shock tube problem at $t = 0.15$ using the compatible energy method and $n_h = 2$ with variable artificial viscosity: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy

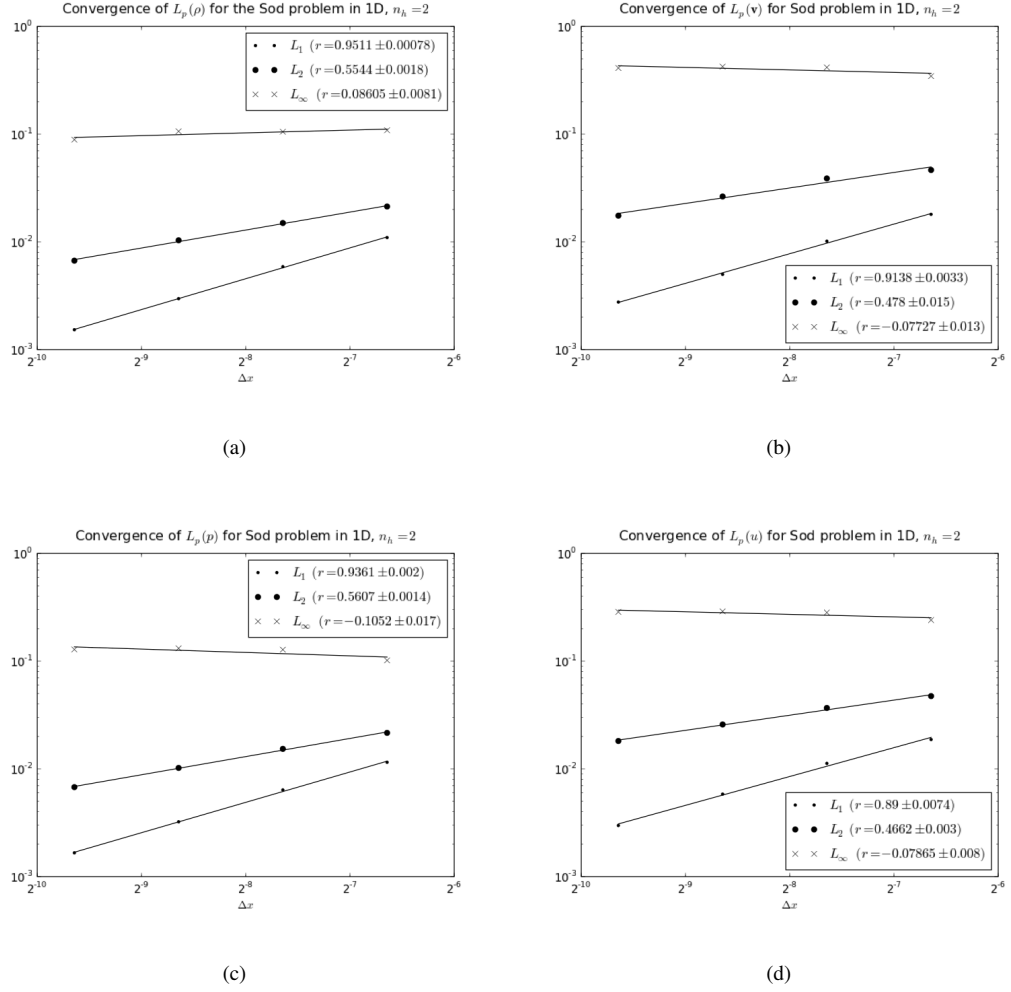


Figure 3.21: The convergence of the Sod shock tube problem at $t = 0.15$ using the compatible energy method and $n_h = 2$ with variable artificial viscosity: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy

to the cells of Athena's grid using the SPH approximation (3.6) and then compute the L_p error norms using Athena's values as the solution. To ensure that the Godunov method is sufficiently resolved, we use a grid with 3200 cells. The SPH solution is shown (dots) imposed upon Athena's solution (solid lines) in Figure 3.22. Figure 3.23 shows the convergence of these error norms for the compatible energy method with $n_h = 2$.






Figure 3.22: The solution of the Sod shock tube problem at $t = 0.15$ using the compatible energy method and $n_h = 2$, compared with that computed by Athena: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy


These convergence rates are very similar to those that we found when comparing the SPH solution to the analytic solution. This gives us confidence that such code-to-code comparisons can be used to perform




(a)



(b)



(c)



(d)

Figure 3.23: The convergence of the Sod shock tube problem at $t = 0.15$ using the compatible energy method and $n_h = 2$, using Athena's solution as the accepted answer: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy

convergence studies in the absence of analytic solutions. This is extremely important to our quantitative assessment of the SPMHD algorithm in the next chapter, since analytic solutions for even the simplest dynamic MHD problems are rare.

The Sod problem has allowed us to demonstrate the quantitative effects of several techniques we discussed earlier in the chapter. We can now assert that the compatible energy method is a measurable improvement to the treatment of weak shocks. We now proceed to a problem involving a strong shock wave.

3.8.4 The Noh problem

A problem described by Noh [125] illustrates the ability of a shock-capturing method to treat shocks of infinite strength. In the Noh problem, a fluid collapses in on itself, and a shock forms, propagating outward from the region of collapse and separating the in-flowing material from stationary, heated fluid. This problem is of particular interest to us, as we study the collapse of a magnetized plasma in Chapter 5. We analyze this problem collapses in several different geometries: the collapse of a fluid in planar, cylindrical, spherical, and conical configurations. As with the Sod problem, we compare our results to analytic solutions and measure the rate of convergence of our form of SPH.

Planar collapse

In the planar Noh problem, two streams of fluid with unit mass density approach one another from unit-length domains to the left and right of the origin. Their inward velocity has magnitude 1, and they have no thermal energy. At time $t = 0$, they touch each other without having interacted. At $t > 0$, the fluids come into contact, and this collision converts all of the kinetic energy within the region of contact into thermal energy. A strong shock wave propagates outward at $v_s = 1/3$, converting the kinetic energy of the inflowing downstream fluid to thermal energy.

This problem may be represented in full or by either of its halves with appropriate boundary conditions: the results are verifiably the same. For computational expediency we represent the two streams (ideal gases with $\gamma = 5/3$) by that at $x > 0$ and a rigid planar boundary condition at $x = 0$. We use equal mass nodes

with a very small amount of thermal energy (corresponding to a pressure $p_0 \approx 10^{-5}$ to prevent spurious negative specific thermal energies), and set $v_x = -1$ for each of them. We set $n_h = 2$ and use the compatible energy method. The solutions for ρ , v_x , p , and u at $t = 0.3$ are shown in Figure 3.24 (with the computed dots laid on top of lines representing Noh's analytic solution); corresponding convergence rates for $N = 100$, $N = 200$, $N = 400$, and $N = 800$ are shown in Figure 3.25.




Figure 3.24: Solution of the Noh problem at $t = 0.3$ using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy

These convergence rates are very close to first order, suggesting that our variant of SPH is treating this problem very well. Specifically, it means that the compatible energy method treats the conversion of

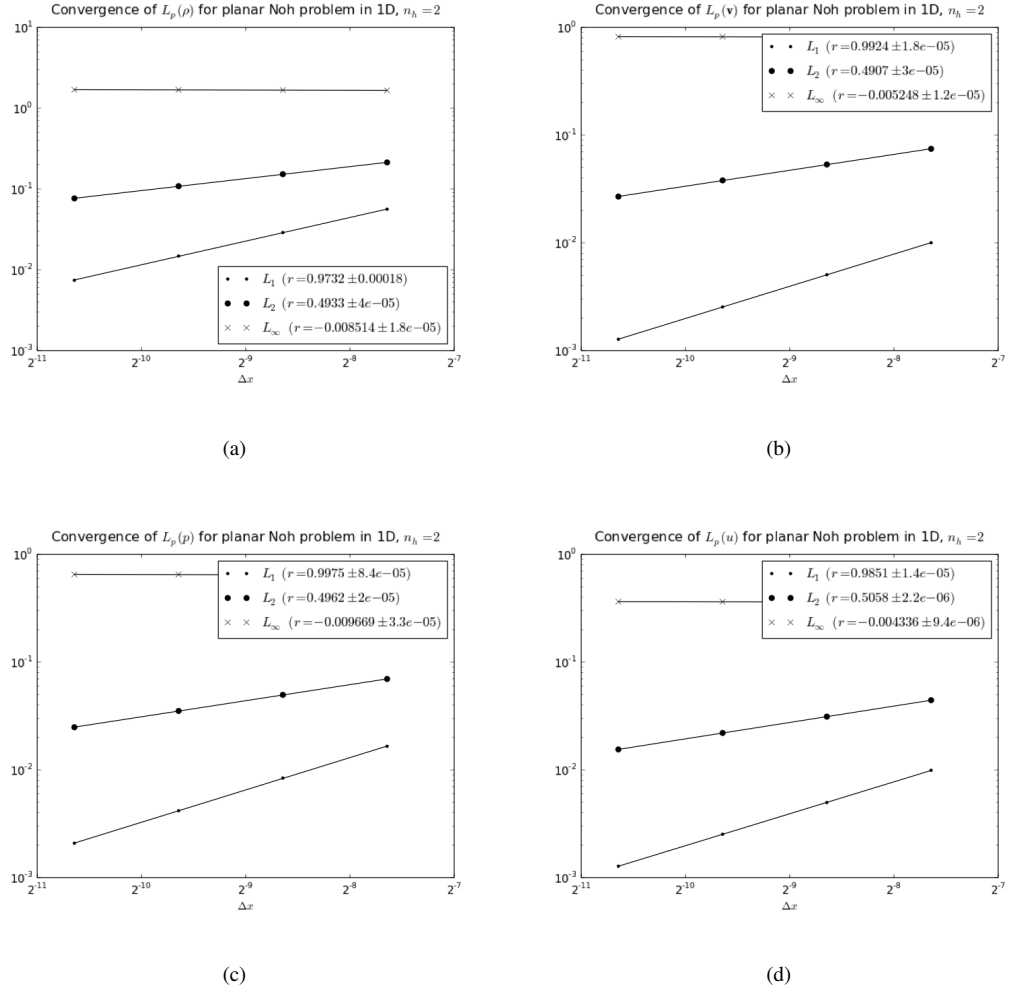


Figure 3.25: Convergence rates for the planar Noh problem using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_x , (c) pressure, and (d) specific thermal energy

kinetic to thermal energy more accurately than the unsymmetric energy equation. These high convergence rates have also been reported by Owen [112]. For comparison, we list the rates of convergence for the L_p norms using the uncorrected energy equation, the ∇h corrections, and the ∇h corrections with the compatible energy method in Table 3.10, Table 3.11, Table 3.12, and Table 3.13.

Method	$r\{L_1(\rho)\}$	$r\{L_2(\rho)\}$	$r\{L_\infty(\rho)\}$
Uncorrected	0.804 ± 0.018	0.393 ± 0.012	-0.03358 ± 0.00078
∇h terms	0.8223 ± 0.014	0.4003 ± 0.01	-0.02605 ± 0.00039
Compatible	0.9732 ± 0.00018	$0.4933 \pm 4e-05$	$-0.008514 \pm 1.8e-05$

Table 3.10: Convergence rates for L_p norms of the mass density in the planar Noh problem using the various energy schemes.

Method	$r\{L_1(\mathbf{v})\}$	$r\{L_2(\mathbf{v})\}$	$r\{L_\infty(\mathbf{v})\}$
Uncorrected	0.8415 ± 0.028	0.3949 ± 0.012	$-0.003253 \pm 1e-06$
∇h terms	0.8388 ± 0.027	0.4026 ± 0.0099	$-4.33e-05 \pm 3.4e-07$
Compatible	$0.9924 \pm 1.8e-05$	$0.4907 \pm 3e-05$	$-0.005248 \pm 1.e-05$

Table 3.11: Convergence rates for L_p norms of the velocity in the planar Noh problem using the various energy schemes.

Method	$r\{L_1(u)\}$	$r\{L_2(u)\}$	$r\{L_\infty(u)\}$
Uncorrected	0.8263 ± 0.014	0.4334 ± 0.005	$0.0005303 \pm 1.5e-8$
∇h terms	0.8346 ± 0.012	0.439 ± 0.0038	$-0.004632 \pm 2.7e-6$
Compatible	$0.9851 \pm 1.4e-5$	$0.5058 \pm 2.2e-6$	$-0.004336 \pm 9.4e-6$

Table 3.12: Convergence rates for L_p norms of the specific thermal energy in the planar Noh problem using the various energy schemes.

When ∇h corrections are used, the convergence rates are improved by very modest amounts, typically around 1%. The compatible energy method is the clear winner in this case. Encouraged by the 1D results, we proceed to study more interesting geometries.

Cylindrical collapse

In the cylindrical case, an ideal gas ($\gamma = 5/3$) with an inward velocity $\mathbf{v} = (v_r, v_\theta, v_z) = (-1, 0, 0)$, density $\rho = 1$, and thermal energy $u = 0$ collapses into the z axis, producing an outward shock with a radial

Method	$r\{L_1(p)\}$	$r\{L_2(p)\}$	$r\{L_\infty(p)\}$
Uncorrected	0.8275 ± 0.023	0.3541 ± 0.023	-0.08796 ± 0.007
∇h terms	0.8571 ± 0.014	0.3736 ± 0.017	-0.07816 ± 0.0052
Compatible	$0.9975 \pm 8.4e-5$	$0.4962 \pm 2e-5$	$-0.009669 \pm 3.3e-5$

Table 3.13: Convergence rates for L_p norms of the pressure in the planar Noh problem using the various energy schemes.

velocity $v_s = 1/3$ behind which sits stationary hot gas. We place nodes on an $N \times N$ lattice on $[-1, 1] \otimes [-1, 1]$ and seed their densities and velocities, giving them minute thermal energies, as in the planar case. We do not use inflow boundary conditions at the outer edge of the problem, since the insertion of points there is more complicated than in the planar case. This “oversight” causes a jump in the L_∞ norms of the computed solutions, as we shall see. These solutions for ρ , v_x , p , and u at $t = 0.3$ using the unsymmetric energy equation are shown in Figure 3.26 (again compared with their analytic solutions). Their convergence rates over the resolutions $N = 50$, $N = 100$, $N = 200$, and $N = 400$ are shown in Figure 3.27. The 2D configurations of the cylindrical implosion at $t = 0$ and $t = 0.6$ appear in Figure 3.28.

These convergence rates are in line with our 1D results. Since there is no shear flow in this problem, we do not observe excess artificial dissipation even with our relatively crude shock capturing scheme. We compare the various corrections we have discussed in the tables Table 3.14 - Table 3.17.

Method	$r\{L_1(\rho)\}$	$r\{L_2(\rho)\}$	$r\{L_\infty(\rho)\}$
Uncorrected	1.016 ± 0.0001	$0.5611 \pm 2.6e-5$	-0.1048 ± 0.0044
∇h terms	0.8122 ± 0.00077	0.4416 ± 0.00018	-0.09473 ± 0.0017
Compatible	1.187 ± 0.0084	0.7234 ± 0.0051	-0.05656 ± 0.0055
$\nabla h + \text{Compatible}$	0.8765 ± 0.0013	0.4917 ± 0.00096	-0.05991 ± 0.00038

Table 3.14: Convergence rates for L_p norms of the mass density in the cylindrical Noh problem using the various energy schemes.

Clearly, the rates of convergence for the compatible energy method are an improvement upon the basic method. Here, as in the 1D results, we observe that the ∇h corrections do not improve the solution beyond that of the compatible energy method. In fact, the solution is worse than the uncorrected solution for $n_h = 2$. The profiles of the solution with these corrections, shown in Figure 3.29, are certainly noisier than




Figure 3.26: Solutions to the cylindrical Noh problem at $t = 0.6$ using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy

Method	$r\{L_1(\mathbf{v})\}$	$r\{L_2(\mathbf{v})\}$	$r\{L_\infty(\mathbf{v})\}$
Uncorrected	1.087 ± 0.0025	0.5742 ± 0.00023	-0.01826 ± 0.00014
∇h terms	0.9285 ± 0.00017	$0.4698 \pm 1.8e-6$	-0.01649 ± 0.00012
Compatible	1.228 ± 0.0024	0.6825 ± 0.0043	-0.0008675 ± 0.00051
∇h + Compatible	0.942 ± 0.0027	0.4921 ± 0.001	$-0.04248 \pm 2e-05$

Table 3.15: Convergence rates for L_p norms of the velocity in the cylindrical Noh problem using the various energy schemes.

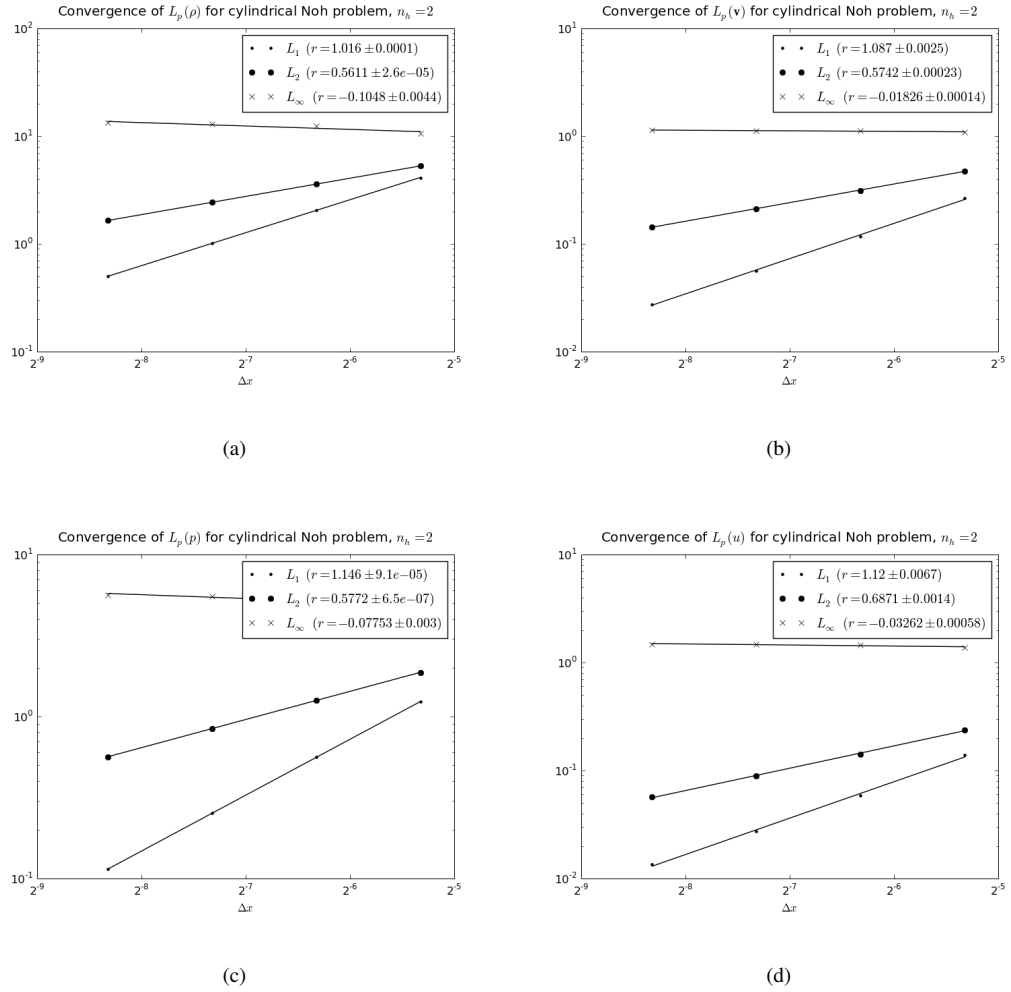



Figure 3.27: Convergence rates for the cylindrical Noh problem using the unsymmetric energy equation with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy


Method	$r\{L_1(u)\}$	$r\{L_2(u)\}$	$r\{L_\infty(u)\}$
Uncorrected	1.12 ± 0.0067	0.6871 ± 0.0014	-0.03262 ± 0.00058
∇h terms	0.9537 ± 0.0019	0.5494 ± 0.0011	-0.03167 ± 0.00047
Compatible	1.256 ± 0.0011	0.7061 ± 0.0024	-0.01371 ± 0.00074
∇h + Compatible	0.9993 ± 0.0048	0.5178 ± 0.0021	-0.01708 ± 0.00018

Table 3.16: Convergence rates for L_p norms of the specific thermal energy in the cylindrical Noh problem using the various energy schemes.




user: jeff
Mon Jan 12 22:41:55 2009

(a)




user: jeff
Mon Jan 12 22:46:49 2009

(b)



user: jeff
Mon Jan 12 22:43:16 2009

(c)



user: jeff
Mon Jan 12 22:43:39 2009

(d)

Figure 3.28: Plots of nodes in the cylindrical Noh problem at $t = 0.6$ using the unsymmetric energy equation with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy

Method	$r\{L_1(p)\}$	$r\{L_2(p)\}$	$r\{L_\infty(p)\}$
Uncorrected	$1.146 \pm 9.1e-5$	$0.5772 \pm 6.5e-7$	-0.07753 ± 0.003
∇h terms	0.8482 ± 0.00061	0.4133 ± 0.00041	-0.08343 ± 0.0022
Compatible	1.244 ± 0.015	0.6974 ± 0.0099	-0.02561 ± 0.0048
$\nabla h + \text{Compatible}$	0.8998 ± 0.0015	0.4589 ± 0.00091	-0.07455 ± 0.0004

Table 3.17: Convergence rates for L_p norms of the pressure in the cylindrical Noh problem using the various energy schemes.

those obtained using the compatible energy method. We have observed that these corrections do improve the solution more at lower values of n_h ; when circumstances such as a tensile instability limit the stencil size, these corrections can be used.

Spherical collapse

The spherical case is similar to the cylindrical case, but with a spherical inward velocity $\mathbf{v} = (v_r, v_\phi, v_\theta) = (-1, 0, 0)$. We can compute the solution to this problem using either the full geometry (an $N \times N \times N$ lattice of nodes on $[-1, 1] \otimes [-1, 1] \otimes [-1, 1]$) or an octant (the lattice on $[0, 1] \otimes [0, 1] \otimes [0, 1]$ with rigid boundary conditions at $x = 0$, $y = 0$, and $z = 0$). Taking advantage of this symmetry, we use our usual parameters (with $n_h = 2$) and compute the solution at resolutions of 10, 20, 40, and 60 nodes on a side. Admittedly, this is very low resolution, but we should still see improvement in the solution as nodes are added. A more thorough study must be conducted to ensure that the convergence does not plateau at higher resolutions. The compatible energy method achieves the best solutions by far for all the methods with the biggest gain in the thermal energy itself; these results are shown in Figure 3.30, and the convergence is plotted in Figure 3.31.

Table 3.18 - Table 3.21 show the convergence rates of the L_p norms using the various methods. The ∇h corrections appear to worsen the solution in this problem, though this could result from the low resolution at which we have run the simulations. Again, we have not imposed inflow boundary conditions at the edge, so the maximum errors reflected in the L_∞ norms are higher than in the planar case.

Once again, the compatible energy method demonstrates real improvements to the modeling of kinetic-to-




Figure 3.29: Solutions to the cylindrical Noh problem at $t = 0.6$ using the compatible energy method and ∇h corrections with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy

Method	$r\{L_1(\rho)\}$	$r\{L_2(\rho)\}$	$r\{L_\infty(\rho)\}$
Uncorrected	0.5874 ± 0.0078	0.3274 ± 0.0058	-0.359 ± 0.028
∇h terms	0.3567 ± 0.017	0.1882 ± 0.0026	-0.2426 ± 0.081
Compatible	0.6216 ± 0.04	0.3835 ± 0.0099	-0.1382 ± 0.016
∇h + Compatible	0.333 ± 0.014	0.2183 ± 0.0024	-0.0261 ± 0.0025

Table 3.18: Convergence rates for L_p norms of the mass density in the spherical Noh problem using the various energy schemes.

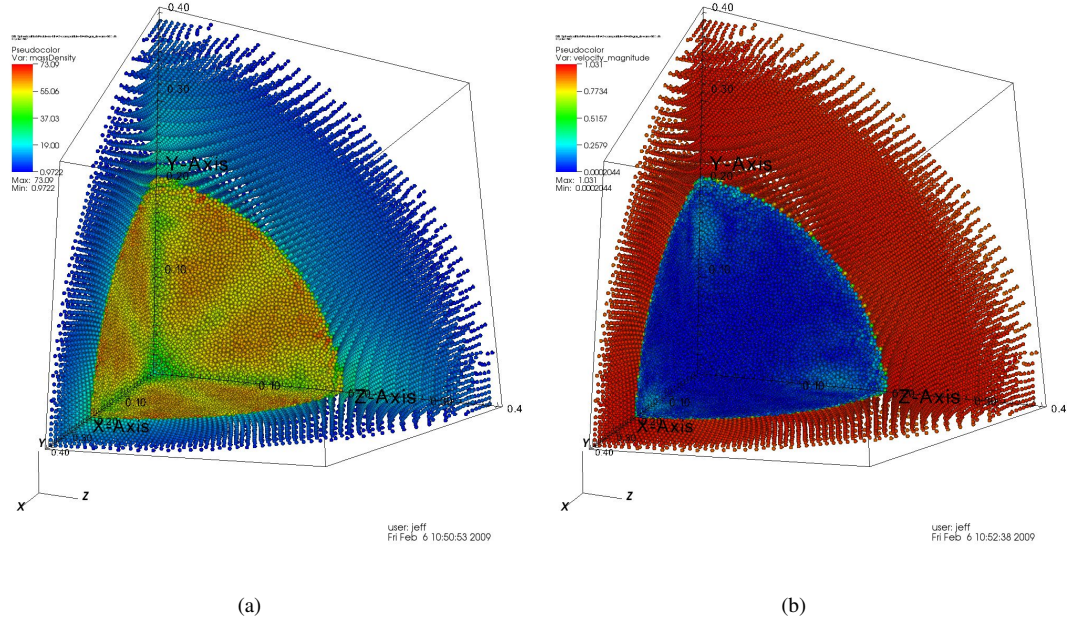


Figure 3.30: Plots of nodes in the spherical Noh problem at $t = 0.6$ using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_r .

Method	$r\{L_1(\mathbf{v})\}$	$r\{L_2(\mathbf{v})\}$	$r\{L_\infty(\mathbf{v})\}$
Uncorrected	0.8664 ± 0.022	0.5245 ± 0.0032	-0.221 ± 0.0021
∇h terms	0.6666 ± 0.0056	0.3596 ± 0.00033	-0.1428 ± 0.0006
Compatible	0.9512 ± 0.00013	0.5172 ± 0.0006	-0.2085 ± 0.0066
∇h + Compatible	0.8261 ± 0.0012	0.4051 ± 0.00047	-0.1577 ± 0.011

Table 3.19: Convergence rates for L_p norms of the velocity in the spherical Noh problem using the various energy schemes.

Method	$r\{L_1(u)\}$	$r\{L_2(u)\}$	$r\{L_\infty(u)\}$
Uncorrected	0.4433 ± 0.054	0.1259 ± 0.023	-0.5971 ± 0.0088
∇h terms	0.4039 ± 0.027	0.1446 ± 0.0092	-0.3677 ± 0.0092
Compatible	1.026 ± 0.00056	$0.5833 \pm 7.3e-5$	-0.07631 ± 0.0048
∇h + Compatible	0.7791 ± 0.0047	0.4151 ± 0.0017	-0.082 ± 0.0021

Table 3.20: Convergence rates for L_p norms of the specific thermal energy in the spherical Noh problem using the various energy schemes.

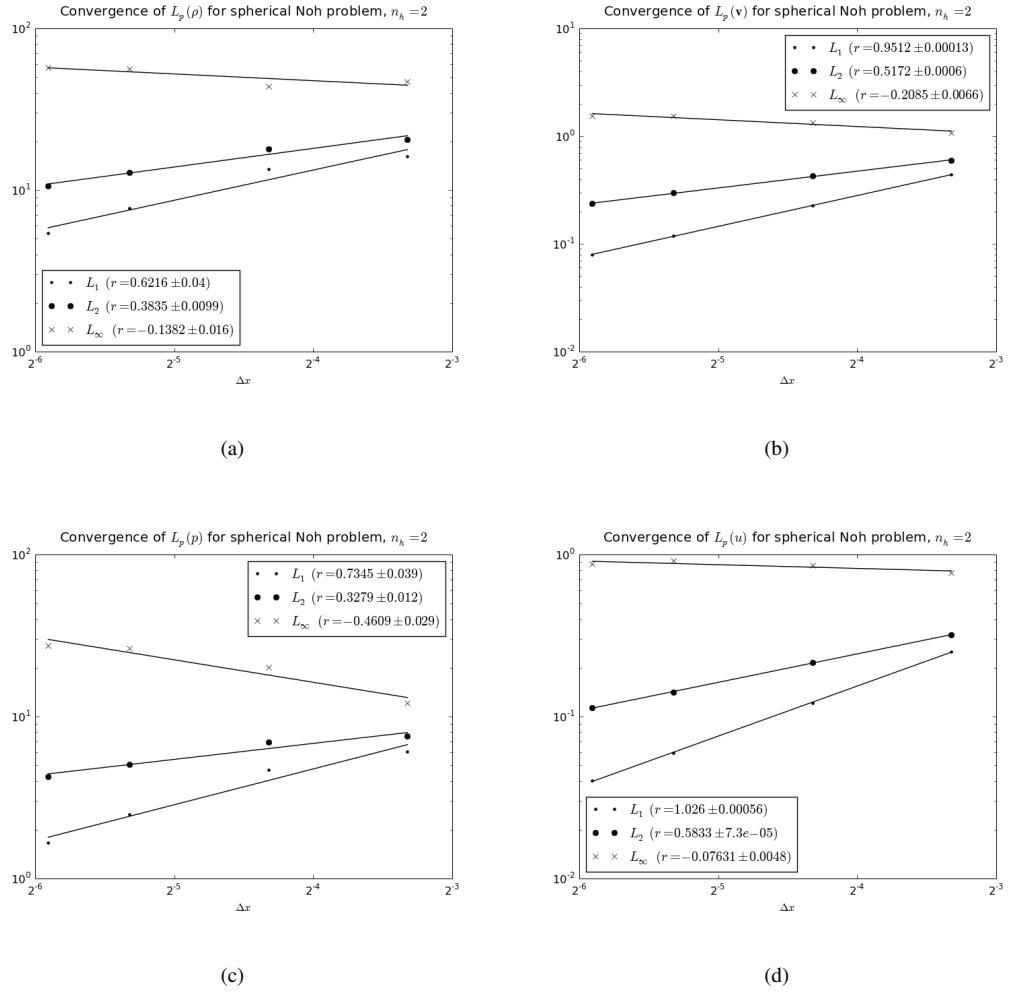


Figure 3.31: Convergence rates for the spherical Noh problem using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy

Method	$r\{L_1(p)\}$	$r\{L_2(p)\}$	$r\{L_\infty(p)\}$
Uncorrected	0.3698 ± 0.00087	0.09609 ± 0.00077	-0.9731 ± 0.035
∇h terms	0.2138 ± 0.00092	0.01665 ± 0.0017	-0.9831 ± 0.13
Compatible	0.7345 ± 0.039	0.3279 ± 0.012	-0.4609 ± 0.029
∇h + Compatible	0.3955 ± 0.0083	0.1499 ± 0.0033	-0.3833 ± 0.014

Table 3.21: Convergence rates for L_p norms of the pressure in the spherical Noh problem using the various energy schemes.

thermal energy conversion and produces good results in three dimensions with no modification.

Conical collapse

The spherical case can be represented as a conical implosion if a rigid boundary condition is used to represent the fluid around a conical region. Here we test a conical boundary condition to make sure it can represent the portion of a sphere outside of a cone of fluid. The nodes are set up on a uniform lattice within this cone, which has unit length and radius. The tip of the cone is at the origin, which is the center of the equivalent sphere. This problem is identical to the spherical Noh problem in every other way.

The method we have used to represent the conical boundary condition is admittedly rather crude, and this problem is a very stringent test of how well it works, especially in representing dynamics near the tip of the cone. In the implosion, the fluid reaches the origin and is stopped as it collides with itself, forming a shock. The boundary nodes created by the conical boundary condition itself do not provide enough force to stop the incoming interior nodes. This is indicated by the “leakage” of the nodes in the vicinity of the origin shown in Figure 3.32. The simulation of this implosion fails to complete because the nodes near the origin are knocked ever further inward by their incoming neighbors, and the post-shock region never stabilizes.

This is not entirely surprising, since the tip of a cone is a complex geometrical object, and we have not programmed the conical surface with any special logic to add extra boundary nodes near the tip. Fortunately, this leakage problem can be fixed by augmenting the conical boundary condition with a rigid planar boundary at $z = 0$. With this additional boundary condition, the post-shock region stabilizes and the simulations complete. The nodes and their values for ρ , \mathbf{v} , and p at $t = 0.6$ are shown in Figure 3.33.

The convergence rates for the L_p norms of ρ , \mathbf{v} , u , and p for the various treatments of the energy equation are given in Table 3.22 - Table 3.25. They are measured for cones with axial node counts 20, 30, 40, 60, and 80. Evidently the compatible energy method produces the best results, particularly improving the otherwise abysmal solutions to the thermal energy in this geometry. The convergence of these quantities for this technique is shown in Figure 3.34. Considering the simplistic representation of the cone, these rates are surprisingly good.

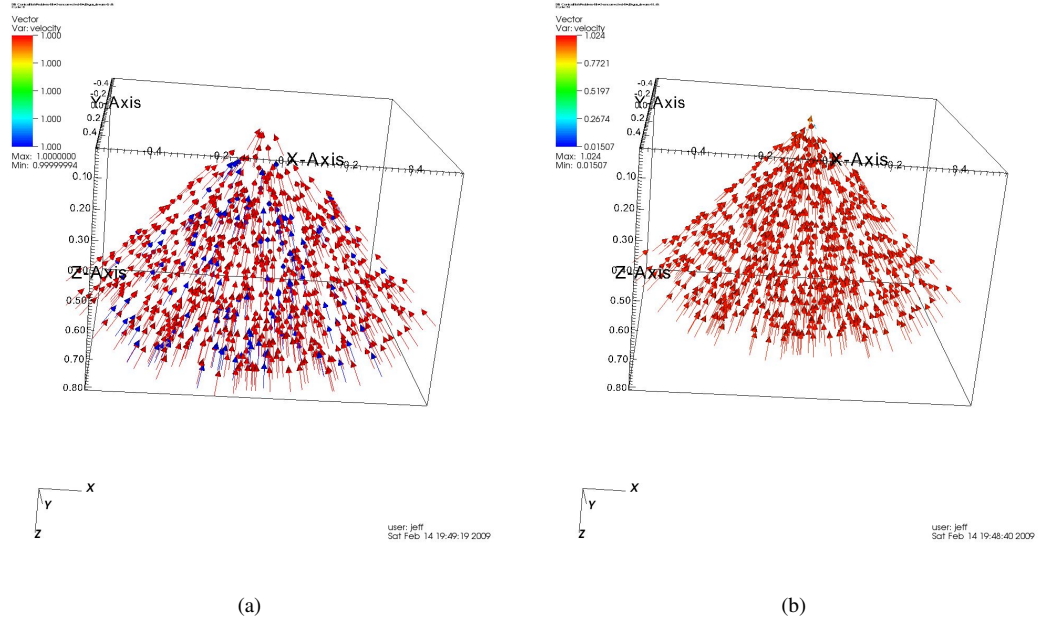


Figure 3.32: (a) Initial velocity configuration of nodes in the conical Noh problem, and (b) leakage of nodes through the origin in the absence of a rigid plane at $z = 0$.

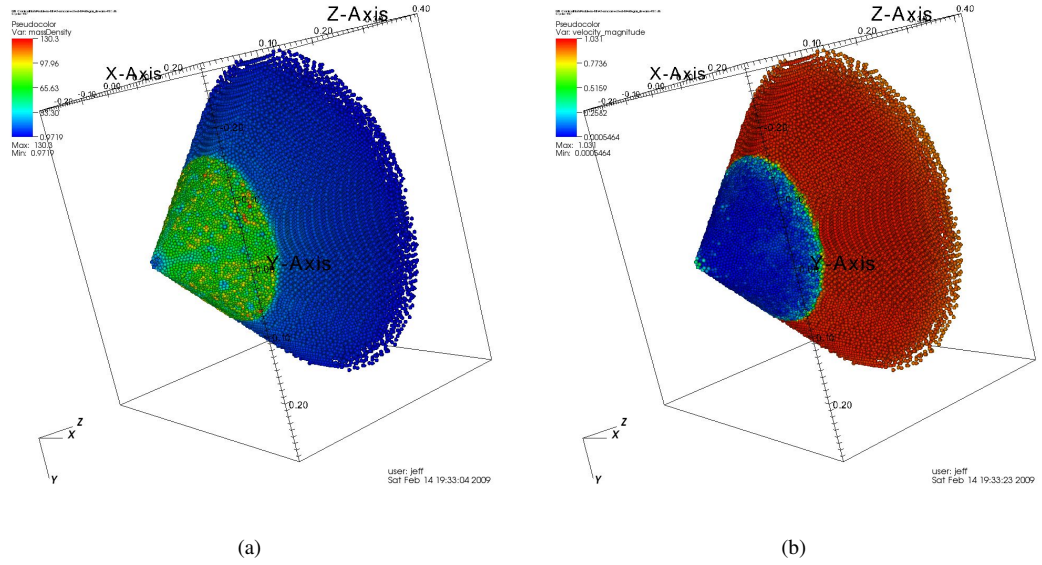


Figure 3.33: Plots of nodes in the conical Noh problem at $t = 0.6$ using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_r .

Method	$r\{L_1(\rho)\}$	$r\{L_2(\rho)\}$	$r\{L_\infty(\rho)\}$
Uncorrected	0.3487 ± 0.048	0.2355 ± 0.013	-0.2878 ± 0.11
∇h terms	0.4047 ± 0.029	0.2024 ± 0.0098	-0.1585 ± 0.021
Compatible	0.6267 ± 0.06	0.3792 ± 0.024	-0.1314 ± 0.02
∇h + Compatible	0.3728 ± 0.017	0.2163 ± 0.0067	-0.1119 ± 0.002

Table 3.22: Convergence rates for L_p norms of the mass density in the conical Noh problem using the various energy schemes.

Method	$r\{L_1(\mathbf{v})\}$	$r\{L_2(\mathbf{v})\}$	$r\{L_\infty(\mathbf{v})\}$
Uncorrected	0.6747 ± 0.033	0.3863 ± 0.0018	-0.1543 ± 0.026
∇h terms	0.5531 ± 0.041	0.2868 ± 0.013	-0.1114 ± 0.015
Compatible	0.7182 ± 0.072	0.4087 ± 0.025	-0.04121 ± 0.004
∇h + Compatible	0.5246 ± 0.032	0.2702 ± 0.0086	-0.08848 ± 0.0075

Table 3.23: Convergence rates for L_p norms of the velocity in the conical Noh problem using the various energy schemes.

Method	$r\{L_1(u)\}$	$r\{L_2(u)\}$	$r\{L_\infty(u)\}$
Uncorrected	0.1145 ± 0.002	-0.04239 ± 0.0026	-0.6163 ± 0.16
∇h terms	0.1894 ± 0.01	0.06186 ± 0.0018	-0.4832 ± 0.088
Compatible	0.7443 ± 0.098	0.4326 ± 0.034	-0.1129 ± 0.015
∇h + Compatible	0.5578 ± 0.045	0.2868 ± 0.012	-0.2198 ± 0.19

Table 3.24: Convergence rates for L_p norms of the specific thermal energy in the conical Noh problem using the various energy schemes.

Method	$r\{L_1(p)\}$	$r\{L_2(p)\}$	$r\{L_\infty(p)\}$
Uncorrected	0.1669 ± 0.0033	0.02568 ± 0.00095	-0.464 ± 0.055
∇h terms	0.1894 ± 0.01	0.06186 ± 0.0018	-0.4832 ± 0.088
Compatible	0.7438 ± 0.077	0.3562 ± 0.022	-0.2013 ± 0.012
∇h + Compatible	0.4026 ± 0.02	0.1687 ± 0.0039	-0.2301 ± 0.023

Table 3.25: Convergence rates for L_p norms of the pressure in the conical Noh problem using the various energy schemes.

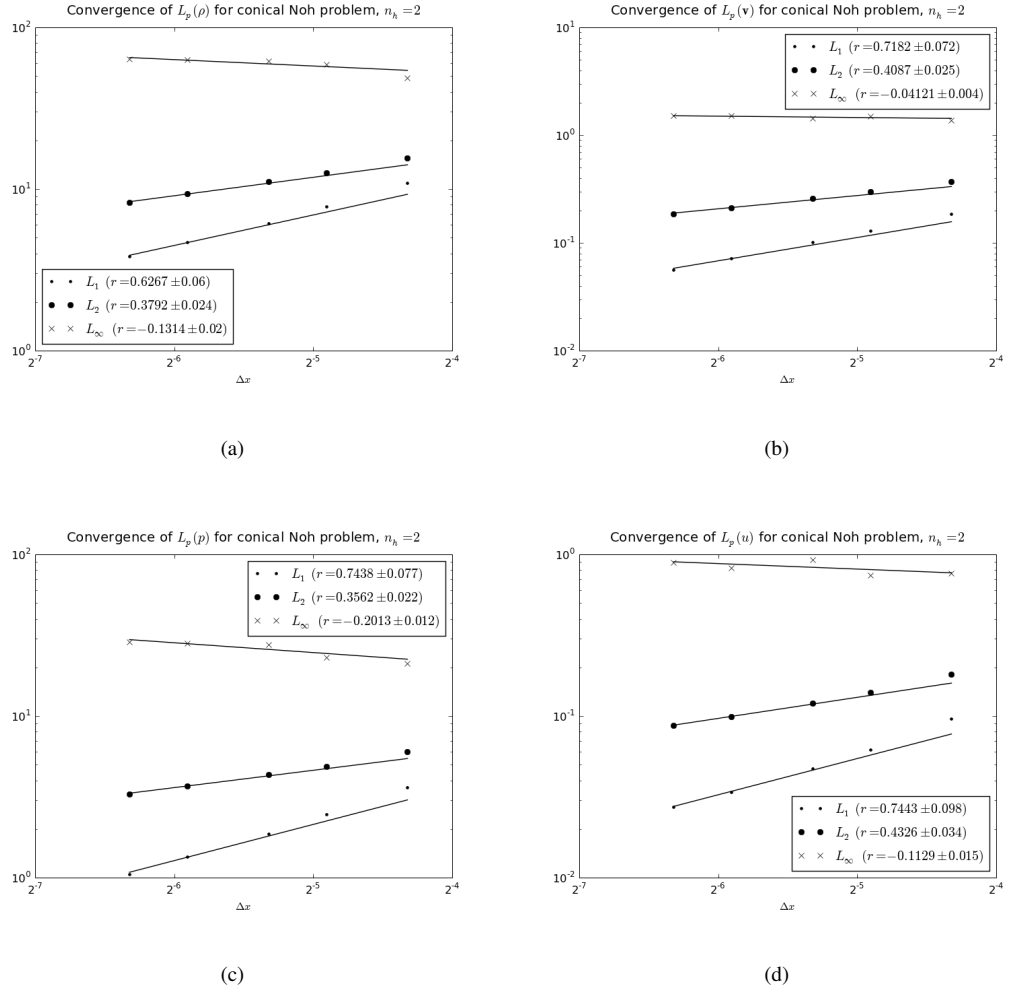


Figure 3.34: Convergence rates for the conical Noh problem using the compatible energy method with $n_h = 2$: (a) mass density, (b) v_r , (c) pressure, and (d) specific thermal energy

The convergence rates in the above tables are lower than those for the spherical Noh problem, but follow similar patterns with respect to the different techniques. We have also used the azimuthal symmetry boundary condition described earlier to model a single quadrant of the cone and have verified that the calculation achieves similar convergence rates. As in the spherical case, the ∇h corrections do not produce good results, and in fact erode the benefits introduced by the compatible energy method when the two methods are combined.

3.9 Summary

In this chapter, we have discussed the important features of SPH and its various parameters. We have discussed some important refinements to the technique, and have shown how these refinements affect the resolving of waves and shocks in ordinary hydrodynamic flows. In particular, we note that the n_h parameter that governs the smoothing scale plays a crucial role in obtaining convergent solutions, and the $n_h = 2$ seems to give optimal results for hydrodynamic simulations. The compatible energy method has produced notable improvements to several of our test problems, and the ∇h corrections that have been discussed by other authors has shown that it is not as helpful as expected when the SPH stencil is increased. The importance of convergence studies in these analyses cannot be overstated—the convergence of the errors in the solutions gives us an objective measure of goodness with which to evaluate these improvements on a problem-by-problem basis. In the next chapter we will add the dynamics of magnetic fields to the method and discuss the ensuing complications.

Chapter 4

Magnetic Fields

We now add magnetic fields to our hydrodynamic equations. In the first section we discuss Smoothed Particle Magnetohydrodynamics as explored by Phillips [1], Børve [2], and Price and Monaghan([98], [90], [84], [3]) and discuss adaptations we have made in pursuit of modeling laboratory plasmas. Since resistivity plays a large role in the behavior of these plasmas, we then describe two ways that the magnetic diffusion equation can be solved without a mesh. We combine these two methods to produce a method for the equations of resistive magnetohydrodynamics. We conclude with test problems that demonstrate the strengths and weaknesses of the method.

4.1 Smoothed Particle Ideal Magnetohydrodynamics

There are myriad formulations of the equations of ideal MHD that use the SPH approximation, but they all share certain features: a magnetic field vector \mathbf{B} is defined on each of the nodes of a conducting medium, and this field exerts a force $\mathbf{J} \times \mathbf{B}$ directly on each of the nodes. Outside of the medium, \mathbf{B} is nominally assumed to be zero. This assumption is both a matter of convenience and a nod to the astrophysical roots of the SPH method. In any case, this thesis focuses on problems in which magnetic fields are present within matter and not in vacuum.

As we have mentioned, the equations of ideal MHD closely resemble those of ordinary hydrodynamics. In fact, the continuity and energy equations are unchanged. We rewrite the momentum equation here, accompanied by the equation for the magnetic field in a perfect conductor:

$$\frac{D\mathbf{v}}{Dt} = -\frac{\nabla p}{\rho} + \frac{\mathbf{J} \times \mathbf{B}}{\rho} \quad (4.1)$$

$$\frac{D\mathbf{B}}{Dt} = \mathbf{B} \nabla \cdot \mathbf{v} - \mathbf{B} \cdot \nabla \mathbf{v} \quad (4.2)$$

It is somewhat discouraging to look at the magnetic induction equation in the Lagrangian form of (4.2). The concept of magnetic flux conservation, a crucial intuitive tool for qualitatively understanding the dynamics of fields within matter, is not expressed very clearly. On the other hand, we have seen that the Lagrangian frame allows us to avoid complications in the momentum equation. Because fields can exist everywhere in space independently of material, neither the Eulerian nor the Lagrangian frame is perfect for solving these equations. A hybrid method that solves the magnetic field would have many advantages, but would come at its own price, which we discuss in Chapter 6.

We will not go into great detail outlining the various SPH discretizations of the MHD equations—their differences are motivated similarly to those of the hydrodynamic equations we discussed in the last chapter. Price gives a more exhaustive survey in his PhD dissertation, associating momentum, energy, and field equations using a discrete SPH-based variational principle [84] in a manner motivated originally by Newcomb [126].

The form of the magnetic induction equation (4.2) chosen by Price is

$$\frac{d\mathbf{B}_i}{dt} = -\frac{1}{\Omega_i \rho_i} \sum_j m_j (\mathbf{v}_{ij} [\mathbf{B}_i \cdot \nabla W_{ij}(h_i)] - \mathbf{B}_i [\mathbf{v}_{ij} \cdot \nabla W_{ij}(h_j)]) \quad (4.3)$$

where we have incorporated the ∇h terms described in the last chapter. These terms may be neglected by setting $\Omega_i = 1$ for each node i . The momentum equation associated with this equation through the variational principle is

$$\frac{d\mathbf{v}_i}{dt} = \sum_j m_j \left[m_j \left(\frac{\mathbf{S}}{\Omega \rho^2} \right)_i \cdot \nabla W_{ij}(h_i) + \left(\frac{\mathbf{S}}{\Omega \rho^2} \right)_j \cdot \nabla W_{ij}(h_j) \right] \quad (4.4)$$

where $\mathbf{S} = -p\mathbf{1} + \frac{1}{\mu_0}\mathbf{B}\mathbf{B} - \frac{1}{2}B^2\mathbf{1}$ is the stress tensor that includes both the hydrodynamic and magnetic forces, and $\mathbf{1}$ is the 3×3 identity tensor. Previously explored algorithms that used approximations of $\mathbf{J} \times \mathbf{B}$ have been found to have poor momentum conservation.

Price describes evolution equations for both the total specific energy $e = v^2/2 + u$ and the specific thermal energy u , but we have found that the results for evolving both of these differ very little, so we will use the same energy equations we used in the last chapter.

4.1.1 The $\nabla \cdot \mathbf{B} = 0$ constraint

As we have mentioned, the magnetic field has an important topological property that is not manifest in the solution of the MHD equations: its divergence is zero. This is of great physical relevance, since the presence of a divergence in \mathbf{B} implies the existence of magnetic monopoles, which have never been observed in nature. Practically speaking, there are three options to those interested in accurately treating MHD:

- One can hope that the discretization of \mathbf{B} is “good enough” to provide relatively divergence-free solutions at some resolution. It is not clear how many scientists follow this approach, but they do so at their own peril: Tóth [26] outlines the hazards of ignoring $\nabla \cdot \mathbf{B}$, not the least of which are large non-physical $\mathbf{J} \times \mathbf{B}$ forces.
- One can design a discretization of \mathbf{B} that renders it manifestly divergence-free. This is the cleanest and most successful approach when it can be employed, but it requires a technique that is either specialized for this purpose or that has some flexibility in the discretization of different quantities. Examples of methods that can represent topological features such as $\nabla \cdot \mathbf{B} = 0$ are the specially-designed Flux-Constrained Transport (FCT) method [33] and compatible methods in the electromagnetics community such as Yee’s Finite Difference Time Domain (FDTD) method [34], the Mimetic Finite Difference (MFD) method [109], and the Vector Finite Element (VFE) Method [38].

- One can periodically “stamp out” the divergence in \mathbf{B} using some sort of projection method. This is a fairly common solution, used both in incompressible hydrodynamic flows([127]) and in MHD([73]). In particular, it allows the use of techniques that were not originally designed to treat magnetic fields. Because of the lack of control we have over the topological structure of nodes in mesh-free methods, we adopt this approach.

In a paper on multidimensional SPMHD, Price outlines a number of candidate methods for divergence cleaning [3]. We have implemented two of these approaches. The first is a modification of a very early approach to SPMHD by Gingold and Monaghan [43] in which \mathbf{B} is projected to the curl of a vector potential by the Biot-Savart law [62]:

$$\mathbf{B} = -\frac{\mu_0}{4\pi} \int \frac{\mathbf{J}(\mathbf{r}') \times (\mathbf{r} - \mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|^3} d\Omega' \quad (4.5)$$

In SPH this projection is performed by computing the current density as the curl of the uncorrected magnetic field \mathbf{B}^* :

$$\mathbf{B} = -\sum_j \frac{m_j \langle \nabla \times \mathbf{B}^* \rangle_j}{4\pi\rho_j} \times \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|^3} \quad (4.6)$$

We use the density-weighted expression (3.25) for $\nabla \times \mathbf{B}^*$. This method removes the divergence from the magnetic field, but suffers from the fact that there are measurable contributions from nodes at large distances. Both Price [3] and Ziegler [128] mention that simulations of multidimensional shock tubes that use periodic boundaries must calculate (4.6) by summing over an infinite number of nodes!

The second method subtracts the gradient of a “damping field” ψ from the induction equation so that the portion of \mathbf{B} having a divergence is damped away. ψ is evolved by the differential equation

$$\frac{d\psi_i}{dt} = -c_i^2 \langle \nabla \cdot \mathbf{B} \rangle_i - \frac{\psi_i}{\tau_i} \quad (4.7)$$

with the signal velocity and time scale defined respectively as

$$c_i^2 = \frac{\gamma P_i}{\rho_i} + \frac{B_i^2}{2\mu_0 \rho_i} \quad \text{and} \quad \tau_i = \frac{n_d h_i}{c_i} \quad (4.8)$$

where n_d is the damping scale length of the divergence as measured in smoothing scales. This method bounds the divergence errors in \mathbf{B} , though it typically does not rid the field of them outright. Nevertheless, we adopt this approach for enforcing the $\nabla \cdot \mathbf{B} = 0$ constraint.

Incidentally, Price mentions in his thesis that the single most significant factor in the divergence error of an SPH node is the number of its neighbors. Since our value of n_h is significantly higher than that of other SPH practitioners, we expect lower divergence errors.

4.1.2 The Tensile Instability

The tension force in (2.54) can excite the tensile instability in SPH as nodes are drawn together to straighten out \mathbf{B} field lines. This instability can be even worse in MHD than in solid mechanics, since $\mathbf{J} \times \mathbf{B}$ forces are often quite strong ($\mathbf{J} \propto \mu_0^{-1} = 10^6/4\pi$). There are a few simple ways of tempering this instability in SPMHD. Børve's approach removes the spurious force along the line-of-sight between nodes. Systems in the presence of a constant external magnetic field \mathbf{B}_0 can be stabilized with a simpler approach: the effects of \mathbf{B}_0 are subtracted from the nodal Maxwell Stress Tensor:

$$\mathbf{T}_i \rightarrow \mathbf{T}_i - \frac{1}{\mu_0} \mathbf{B}_0 \mathbf{B}_0 \quad (4.9)$$

In the continuum case, this should have no effect, since a constant vector field should not contribute to the magnetic stress. However, the differencing of $\nabla \cdot \mathbf{T}$ produces nonzero results even in the case of a constant field, a consequence of the *inconsistency* of SPH. In 1D, we can use this method to ignore the magnetic stress caused by B_x , since $\nabla \cdot \mathbf{B} = 0$ implies that $\partial B_x / \partial x = 0$ in that case. Phillips [1] used a similar method that subtracts the maximum nodal stress in the domain from the stress tensor at each node.

Morris [47] uses an altogether-different approach, sacrificing the conservation of momentum but using a discrete form of the gradient that is exactly zero in the presence of constant magnetic fields. We do not

explore this technique because momentum conservation is very important the treatment of shock problems. We do note this is a step in the direction of a consistent SPH method such as Corrected SPH (CSPH) [80], a method that we advocate in the future treatment of SPH in the presence of magnetic fields.

Even with these methods, our test problems show that the tensile instability is significant enough to cause trouble to the extent that we must revise our evaluation of the n_h parameter. In particular, we find that setting n_h to 2 (a good rule of thumb for pure hydrodynamic flows) often leads to oscillations that can be quelled by setting n_h to a lower value. We find that both the presence of these oscillations and the optimal value for n_h is quite problem dependent, though.

4.1.3 Smoothing field discontinuities

Because magnetic field discontinuities beget discontinuities in the corresponding transverse components of the fluid velocity, Price has developed an artificial dissipation model for magnetic fields in ideal MHD [98]. The form of the dissipation is a dissipative term in the induction equation

$$\frac{d\mathbf{B}_i}{dt} = \rho_i \sum_j m_j \frac{\alpha_B v_{sig}}{2\bar{\rho}_{ij}^2} (\mathbf{B}_i - \mathbf{B}_j) \hat{\mathbf{r}}_{ij} \cdot \nabla W_{ij} \quad (4.10)$$

and mimics the presence of a magnetic diffusion term $-\nabla \times (\eta \nabla \times \mathbf{B})$ with a value for η that depends essentially on α_B . This adds a corresponding Joule heating term

$$\frac{du_i}{dt} = - \sum_j m_j \frac{v_{sig}}{2\bar{\rho}_{ij}} \left[\frac{\alpha_B}{2\mu_0 \bar{\rho}_{ij}} |\mathbf{B}_i - \mathbf{B}_j|^2 \right] \hat{\mathbf{r}}_{ij} \cdot \nabla W_{ij}. \quad (4.11)$$

In his work, Price used a variable switch α_B that is evolved using an equation very similar to that used by Morris for artificial viscosity, but the resulting oscillations in \mathbf{B} and in \mathbf{v} ultimately led him to peg its value at 1. In laboratory plasmas, the use of such “artificial resistivity” is less crucial, since this term is only necessary for modeling magnetohydrodynamic flows in perfect conductors—adding physical resistivity for laboratory plasmas typically provides sufficient smoothing for \mathbf{B} .

4.1.4 The Compatible Energy Method Extended to MHD

In Chapter 3 we outlined a method developed by Owen [112] for evolving the specific thermal energy that exactly conserves total energy. We have extended this method for magnetohydrodynamic flows by applying the same discrete conservation principle in the presence of magnetic fields. This formulation has some pitfalls that do not appear in the case of ordinary hydrodynamics: for example, the most straightforward way to derive a compatible MHD energy update is to begin with the expression for total hydromagnetic energy:

$$E^n = \sum_i m_i \left(\frac{1}{2} (v_i^n)^2 + u_i^n + \frac{1}{2} \frac{(B_i^n)^2}{\mu_0 \rho_i} \right) \quad (4.12)$$

If we proceed to add the extra terms from the magnetic field into the expressions for the hydrodynamic update and incorporate the Lorentz force into the pairwise accelerations \mathbf{a}_{ij} in (3.90), we observe that there is energy transfer between the magnetic field and the thermal motion that violates thermodynamic principles: in particular, the magnetic field is able to steal energy directly from thermal energy, resulting in a physical “Joule cooling.” Therefore, we must approach the technique with more care.

In ideal MHD, no Joule heating occurs, and no direct transfer of energy is allowed between the magnetic field and the thermal energy. This means that any contribution to the thermal energy from pairwise accelerations must discount any acceleration produced by $\mathbf{J} \times \mathbf{B}$. Since the contribution to the thermal energy of node i from node j is

$$\Delta u_{ij} = f_{ij} \frac{\Delta E_{ij}^{thermal}}{m_i} = f_{ij} \left(\mathbf{v}_j^{n+1/2} - \mathbf{v}_i^{n+1/2} \right) \cdot \mathbf{a}_{ij}^n \Delta t \quad (4.13)$$

we must remove the influence of $\mathbf{J} \times \mathbf{B}$ from $\mathbf{v}_i^{n+1/2}$, $\mathbf{v}_j^{n+1/2}$, and \mathbf{a}_{ij}^n . Thus we transform the half-step velocities:

$$\mathbf{v}_i^{n+1/2} \rightarrow \mathbf{v}_i^{n+1/2} - \frac{\Delta t}{2m_i} (\mathbf{J} \times \mathbf{B})_i^n \quad (4.14)$$

$$\mathbf{v}_j^{n+1/2} \rightarrow \mathbf{v}_j^{n+1/2} - \frac{\Delta t}{2m_j} (\mathbf{J} \times \mathbf{B})_j^n \quad (4.15)$$

Meanwhile, the pairwise accelerations \mathbf{a}_{ij}^n do not include terms from $\mathbf{J} \times \mathbf{B}$. With these modifications, (3.90) may be used to update the thermal energy. In the case of resistive MHD or in cases in which Price's artificial thermal energy term (3.56) is used, the expression for the pairwise change in total energy (3.87) must also be modified:

$$\Delta E_{ij}^{thermal} \rightarrow \Delta E_{ij}^{thermal} + m_i \Delta t \left(\frac{\partial u}{\partial t} \right)_{ij}^{diss} + m_j \Delta t \left(\frac{\partial u}{\partial t} \right)_{ji}^{diss} \quad (4.16)$$

where $(\partial u / \partial t)_{ij}^{diss}$ is the sum of all dissipative sources of heat on node i due to the influence of node j .

In preserving the physics preventing the magnetic field from cooling the nodes, we have sacrificed the exact conservation achieved by the purely-hydrodynamic compatible energy formulation. The inherent difficulty of exact conservation in hydromagnetic energy stems from the fact that \mathbf{v} and \mathbf{B} are evolved by separate differential equations whose discrete representation does not respect the relationship of the continuous velocity to its magnetic field. Thus, any attempt we make to enforce the conservation of the discrete hydromagnetic energy (4.12) pushes all of the discretization error into the specific thermal energy. In this case, obeying the laws of thermodynamics is definitely preferable to achieving exact energy conservation, so this sacrifice is acceptable.

4.1.5 Boundary Conditions on \mathbf{B}

The boundary conditions relevant to magnetic fields for MHD flows are similar in nature to those for hydrodynamic flows, and we find that periodic boundaries, inflow and outflow boundaries, Dirichlet boundaries, and free boundaries can be treated for \mathbf{B} just as they are in ordinary SPH.

Reflected boundary conditions on \mathbf{B} are not typically used, since they give rise to extremely large surface currents at the boundary. A case of interest that is not covered by methods we have discussed is the

behavior of \mathbf{B} near a conducting boundary. The physics of this boundary condition is discussed in Chapter 2. We defer this discussion to the discussion of boundary conditions for the SPH method of treating magnetic diffusion in the following section.

4.2 Magnetic Diffusion in SPH

We have briefly discussed the magnetic diffusion equation (2.35), and its role in smoothing a magnetic field within an imperfect conductor. We obtained this equation from Faraday’s Law (2.6) by relating the electric field to the current density that it induces:

$$\mathbf{E} = \eta \mathbf{J} \quad (4.17)$$

To this point we have treated electrical resistivity as a scalar quantity η . This implies that current flows purely in the direction of the electric field that induces it: \mathbf{J} must be parallel to \mathbf{E} . Sometimes, an electric field can induce currents that are not aligned with it. For example, charge carriers can be scattered by ions in a plasma so that there exists some transverse component of \mathbf{J} . In this case the resistivity of the plasma is anisotropic and is represented by a symmetric tensor \mathbf{R} , viz.

$$\mathbf{E} = \mathbf{R} \mathbf{J} \quad (4.18)$$

Thus, in its most general form, the magnetic diffusion equation is

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times (\mathbf{R} \cdot \nabla \times \mathbf{B}) \quad (4.19)$$

The resistivity tensor is necessarily positive-definite. If it weren’t, then \mathbf{J} could flow in the opposite direction of \mathbf{E} —an amusing proposition. In this case, the plasma would be doing work on the field, exchanging thermal energy for field energy via “Joule cooling” in which $\mathbf{J} \cdot \mathbf{E} < 0$. Note that if η is allowed to vary in space, we must compute its gradient as well.

4.2.1 Laplacian Operators in SPH

There are a number of options for practically discretizing the Laplacian operator. Colin et al [88] use a Difference Gradient Approximation form for the Laplacian based on the Hodge decomposition for incompressible flows—this decomposition can probably be applied to the “vector” Laplacian as well. Unfortunately, it is unclear how this method can accommodate a Neumann boundary condition in any straightforward way. In the case of a scalar field ϕ , it can treat Neumann boundary conditions in the form of a specified normal derivative $\partial\phi/\partial n$ at a boundary, but this technique will not work for specifying the curl of a vector field. This is because the constraint equations for specifying the curl produce rows in the Laplacian matrix that render it non-positive-definite and likely singular unless boundary terms are included.

Remeshed SPH (RSPH) [129] uses periodic remapping to overcome inaccuracies and instabilities associated with the ordinary SPH Laplacian operator. When the nodes become too disordered, they are remapped to a more regular configuration. While this allows one to use the ordinary SPH Laplacian, its complications (the recognition of sufficient node disorder to remap, the remapping process) are a significant price to pay for this convenience. And again, it is not clear how Neumann boundary condition would be applied.

Corrected SPH again offers improvements in this area, with its corrected expressions for gradients and much-improved accuracy at the boundary. This method has been used to construct a form for the Poisson equation using a variational principle [130], conferring some of the advantages of a weak-form approach without incurring costly numerical integration. More recently, it has been used to solve problems in thermal diffusion [131] near free boundaries.

The method we have used based on work by Leigh Brookshaw, in which he derives a form for the Laplacian by using ∇W as an interpolant instead of W . This method has been used successfully to model heat conduction [4] and radiation transport in the diffusion approximation [46]. We consider the Taylor expansion of a magnetic field about a point \mathbf{x}_i , evaluated at a point \mathbf{x}_j :

$$\mathbf{B}(\mathbf{x}_j) = \mathbf{B}(\mathbf{x}_i) + \nabla \mathbf{B}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i) + \frac{1}{2} (\mathbf{x}_j - \mathbf{x}_i) \cdot \nabla \nabla \mathbf{B}(\mathbf{x}_i) \cdot (\mathbf{x}_j - \mathbf{x}_i) + O(\Delta x^3). \quad (4.20)$$

Defining $\mathbf{x}_{ji} \equiv (\mathbf{x}_j - \mathbf{x}_i)$ and $\mathbf{B}_i \equiv \mathbf{B}(\mathbf{x}_i)$ (and the like), we approximate the Taylor expansion with an integral, multiplying (4.20) by the expression

$$\hat{\mathbf{x}}_{ji} \frac{\nabla W(\mathbf{x}_{ji}; h_i)}{|\mathbf{x}_{ji}|^2}. \quad (4.21)$$

and integrating over all \mathbf{x} . The motivation for doing this is that (4.21) is well-behaved as $\mathbf{x}_j \rightarrow \mathbf{x}_i$, making the integral approximation less sensitive to node disorder. The symmetry of the kernel W (and the resulting anti-symmetry of its gradient) ensures that the terms involving $\nabla \mathbf{B}$ and the off-diagonal terms of the Hessian matrix $\nabla \nabla \mathbf{B}$ vanish, and we are left with an expression for $\nabla^2 \mathbf{B}$:

$$\langle \nabla^2 \mathbf{B} \rangle_i = -2 \int \frac{\mathbf{B}_j - \mathbf{B}_i}{|\mathbf{x}_{ij}|^2} \mathbf{x}_{ij} \cdot \nabla W_{ij} d\Omega \quad (4.22)$$

This is sufficient for the representation of a field in a homogeneous resistor, since $-D \nabla \times \nabla \times \mathbf{B} = D \nabla^2 \mathbf{B}$ by (2.33), where we have used the magnetic diffusivity $D = \eta/\mu_0$. If η varies in space, we must use the more general expression

$$-\nabla \times (D \nabla \times \mathbf{B}) = \nabla(D \nabla \mathbf{B}) - \nabla \mathbf{B} \cdot \nabla D. \quad (4.23)$$

One way to express the $\nabla(D \nabla \mathbf{B})$ term is to use the identity

$$\nabla(D \nabla \mathbf{B}) = \frac{1}{2} [\nabla^2(D \mathbf{B}) - \mathbf{B} \nabla^2 D + D \nabla^2 \mathbf{B}]. \quad (4.24)$$

Using (4.22) and (4.24), it is easy to obtain the SPH expression

$$\langle \nabla(D \nabla \mathbf{B}) \rangle_i = \sum_j (D_i + D_j) \frac{m_j}{\rho_j} \frac{\nabla W_{ij}}{|\mathbf{x}_{ij}|^2} \quad (4.25)$$

for the first term in (4.23). This is much more accurate in practice than using the product rule $\nabla(D\nabla\mathbf{B}) = \nabla D \cdot \nabla\mathbf{B} + D\nabla^2\mathbf{B}$ and the SPH gradient approximation for continuous diffusivities. Here, we have simply averaged the diffusivities of nodes i and j . In their treatment of heat transfer, Monaghan and Cleary [45] recommend the substitution

$$\frac{D_i + D_j}{2} \rightarrow D_{ij} = \frac{2D_i D_j}{D_i + D_j} \quad (4.26)$$

which ensures that the magnetic diffusivity always remains continuous. We adopt this recommendation when we use this approach. For the second term in (4.23), we use the density-weighted gradient (3.23) for ∇D .

An explicit treatment of magnetic diffusion would simply traverse the nodes in the distribution and accumulate contributions to the discrete magnetic diffusion equation such as

$$\frac{\partial \mathbf{B}_i}{\partial t} = \sum_j (D_{ij}) \frac{m_j}{\rho_j} \frac{\nabla W_{ij}}{|\mathbf{x}_{ij}|^2} (\mathbf{B}_i - \mathbf{B}_j) - \sum_j \frac{m_j}{\rho_j} \nabla W_{ij} \cdot \mathbf{B}_j \cdot \frac{1}{\rho_j} \sum_k m_k D_k \nabla W_{ij}. \quad (4.27)$$

However, explicit integration methods perform poorly for diffusion problems because they impose severe constraints on the size of the time step. Thus we represent $\partial \mathbf{B} / \partial t$ by a linear operator \mathcal{D} acting on a row vector \mathcal{B} :

$$\frac{\partial \mathcal{B}}{\partial t} = \mathcal{D}(\mathcal{B}) \quad (4.28)$$

This allows us to employ an implicit time integration scheme that we describe shortly.

The diffusion matrix \mathcal{D} for an N -node distribution is a $3N \times 3N$ matrix whose off-diagonal and diagonal elements are given respectively by

$$\begin{pmatrix} \mathcal{D}_{3i,3j} & \mathcal{D}_{3i,3j+1} & \mathcal{D}_{3i,3j+2} \\ \mathcal{D}_{3i+1,3j} & \mathcal{D}_{3i+1,3j+1} & \mathcal{D}_{3i+1,3j+2} \\ \mathcal{D}_{3i+2,3j} & \mathcal{D}_{3i+2,3j+1} & \mathcal{D}_{3i+2,3j+2} \end{pmatrix} = (D_{ij}) \frac{m_j}{\rho_j} \frac{\nabla W_{ij}}{|\mathbf{x}_{ij}|^2} \mathbf{1} - \frac{m_j}{\rho_j^2} \nabla W_{ij} \otimes \sum_k m_k D_k \nabla W_{ij} \quad (4.29)$$

$$\begin{pmatrix} \mathcal{D}_{3i,3i} & 0 & 0 \\ 0 & \mathcal{D}_{3i+1,3i+1} & 0 \\ 0 & 0 & \mathcal{D}_{3i+2,3i+2} \end{pmatrix} = - \sum_j (D_{ij}) \frac{m_j}{\rho_j} \frac{\nabla W_{ij}}{|\mathbf{x}_{ij}|^2} \mathbf{1} \quad (4.30)$$

where $\mathbf{1}$ is the 3×3 identity tensor. Accordingly, \mathcal{B} is a $3N$ -row vector for which

$$\mathcal{B}_{3i} = B_x(\mathbf{x}_i), \quad \mathcal{B}_{3i+1} = B_y(\mathbf{x}_i), \quad \mathcal{B}_{3i+2} = B_z(\mathbf{x}_i) \quad (4.31)$$

It is evident from (4.29) that \mathcal{D} is not a symmetric matrix. Fortunately, this is not an obstacle to obtaining accurate solutions of the magnetic diffusion equation. And since there is no numerical integration involved in the construction of the matrix (in contrast to finite element methods and their ilk), \mathcal{D} can be constructed very quickly even when the node connectivity changes every integration step.

Time Integration: Fully Discrete Form

The characteristic time scale over which diffusion occurs ((2.58)) can serve as a CFL-like relationship between the time step size and the nodal spacing. Unfortunately, if one adopts this approach, one finds quickly that a time step that scales quadratically with the nodal spacing becomes impractically small for realistic problems. To avoid this issue, we adopt a variable, implicit, Crank-Nicholson-based time integration method that has become popular for solving diffusion and diffusion-like equations [[132], [38]]. We begin with our semi-discrete linear system (4.28), using the same superscript/subscript notation we introduced in Chapter 3. We introduce a weighting factor θ (with $0 \leq \theta \leq 1$) to control the “implicitness” of the time discretization:

$$\left(\frac{d\mathcal{B}}{dt} \right)^{n+\frac{1}{2}} = \frac{1}{\mu_0} \mathcal{D}(\eta^{n+\theta}) \mathcal{B}^{n+\theta} \quad (4.32)$$

We then approximate

$$\left(\frac{d\mathcal{B}}{dt}\right)^{n+\frac{1}{2}} \approx \frac{\mathcal{B}^{n+1} - \mathcal{B}^n}{\Delta t} \quad \text{and} \quad \mathcal{B}^{n+\theta} \approx \theta \mathcal{B}^{n+1} + (1-\theta)\mathcal{B}^n \quad (4.33)$$

where $\Delta t = t^{n+1} - t^n$, rearranging the terms to obtain

$$\left[\mathbf{1} - \theta \Delta t \mu_0^{-1} \mathcal{D}(D^{n+\theta})\right] \mathcal{B}^{n+1} = \left[\mathbf{1} + (1-\theta) \Delta t \mu_0^{-1} \mathcal{D}(D^{n+\theta})\right] \mathcal{B}^n \quad (4.34)$$

This linear system must be solved for each time t_n to evolve the magnetic induction to its state at time t^{n+1} .

The time integration scheme is unconditionally stable for $\theta \geq 1/2$. If $\theta = 0$, we recover the explicit Forward Euler integration method. If $\theta = 1$, this discretization gives us the implicit Backward Euler scheme. Both of these are first-order accurate in time. If $\theta = 1/2$, we get the second-order Crank-Nicholson scheme. With this scheme, we can scale the time step size linearly with the nodal spacing and still obtain accurate, stable results.

Note that θ also determines the time centering of the electrical resistivity η in (4.32), which must be computed at $t^{n+\theta} = t^n + \theta \Delta t$. This is crucial if one wishes to obtain the predicted convergence rate for the selected time discretization.

For test problems involving magnetic diffusion, we typically use $\theta = 1/2$ and scale the maximum time step size down with the nodal spacing to control the amount of time discretization error in the solution. If one were to use a value of θ less than $1/2$, one must take care to ensure that the spectral radius of the amplification matrix

$$\mathbf{A} = \left[\mathbf{1} - \theta \Delta t \mu_0^{-1} \mathcal{D}(D^n)\right]^{-1} \left[\mathbf{1} + (1-\theta) \Delta t \mu_0^{-1} \mathcal{D}(D^n)\right] \quad (4.35)$$

is less than 1 in order to achieve stability [101]. Alternatively, one may use (2.58) as an estimate for the maximum time step (though, as previously mentioned, this becomes impractical as the nodal spacing in a problem becomes small). In practice, one must be careful even when $\theta = 1/2$, as problems with discontinuities may develop oscillations; these can be eliminated by increasing θ or more carefully controlling the time step size.

We solve (4.34) using an iterative Generalized Minimum Residual (GMRES) solver with a residual error tolerance of 10^{-18} . This method does not require that the matrix in our system be symmetric. We use a simple Jacobi preconditioner, since that is sufficient to obtain convergent solutions with this method.

Once the new value of \mathbf{B} is obtained in this manner, we compute the current density from (2.27) using the mass-density-weighted form of the SPH curl operator (3.25).

4.2.2 Boundary Conditions

The magnetohydrodynamic equations we discussed in the previous chapter propagate signals at several characteristic speeds (the fast and slow magnetosonic speeds, the Alfven speed) across a domain. In this sense, causality is encoded within these equations: we may be assured that nothing happens hydrodynamically in a region that a signal has not reached. In diffusion equations, there is no such characteristic speed: the entire domain responds instantaneously to a diffusing quantity. Strictly, this is not physically correct, since changes in \mathbf{B} can move no faster than the speed of light. But our choice to ignore the time scale of electron motion (and everything quicker) means that the field is allowed to behave this way.

Dirichlet Boundaries

The greatest consequence of this “acausality” is that we must be prepared to directly specify the behavior of \mathbf{B} or its derivative along the entire boundary of the domain, Γ . If we know the value of the field itself on a portion of the boundary which we denote $\Gamma_{\mathbf{B}}$, we may (borrowing the language of finite elements) specify an *essential* boundary condition on \mathbf{B} :

$$\mathbf{B}(\mathbf{x}, t) = \mathbf{B}_0(\mathbf{x}, t) \text{ on } \Gamma_{\mathbf{B}} \quad (4.36)$$

As in the hydrodynamic case, we set up ghost nodes on the other side of the Dirichlet boundary and ensure that their \mathbf{B} values are set properly. This requires that we have a function $\mathbf{B}(\mathbf{x}, t)$ defined outside the boundary, but it is often not difficult to make that extension.

\mathbf{B} may also be specified by a surface current and a normal magnetic field on another portion of the boundary $\Gamma_{\mathbf{K}}$, as in (2.43):

$$\mathbf{B} = B_n \mathbf{n} - \mathbf{n} \times \mathbf{K} \text{ on } \Gamma_{\mathbf{K}} \quad (4.37)$$

In this case, we want \mathbf{B} to have its tangential and normal components near the boundary—we don't actually model the field *at* or *beyond* the boundary. The most straightforward way to mock up this configuration is to evaluate $\mathbf{B}_{\Gamma_{\mathbf{K}}} = B_n(\mathbf{x}, t) \mathbf{n} - \mathbf{n} \times \mathbf{K}(\mathbf{x}, t)$ for a set of projected points $\{\mathbf{x}\}$ on $\Gamma_{\mathbf{K}}$ and construct ghost nodes as we did with the previous Dirichlet boundary, assigning $\mathbf{B}_{\Gamma_{\mathbf{K}}}$ to each of these nodes.

Neumann Boundaries

On another part of the boundary $\Gamma_{\mathbf{J}}$, we may not know the behavior of \mathbf{B} itself, but instead may wish to specify that of \mathbf{J} in the form of a *natural* boundary condition:

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} \text{ on } \Gamma_{\mathbf{J}} \quad (4.38)$$

We can use this boundary condition to represent a plasma-vacuum interface at which $\mathbf{J} = \nabla \times \mathbf{B} = \mathbf{0}$ because of the lack of charge carriers in vacuum. Conveniently, the Brookshaw method has this boundary condition built into it at free boundaries, since the lack of neighboring nodes $\{j\}$ at the boundary gives no contribution to the Laplacian (4.22). Unfortunately, specifying an arbitrary current \mathbf{J} at the boundary is more difficult, and we have not determined how to do so accurately with this method. Were we to model problems with known currents or electric fields (such as those in plasma acceleration), for which

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{R}^{-1} \mathbf{E} \text{ on } \Gamma_{\mathbf{E}}, \quad (4.39)$$

we would probably use another method for solving the magnetic diffusion equation. We explore such an alternative in the next section. CSPH, with its accuracy near the boundary, also should be explored, but it is beyond the scope of this thesis to do so.

4.2.3 Limitations

The simplicity and robustness of Brookshaw’s method limits its flexibility in a few ways. Because of the way we represent the $\nabla \times \nabla \times$ operator directly as the Laplacian ∇^2 , there is no obvious way to use the method to model diffusion with an anisotropic resistivity \mathbf{R} . Also, we have mentioned how the method automatically satisfies the homogeneous Neumann boundary condition representing $\mathbf{J} = \mathbf{0}$ at the surface of a conductor, but that in the presence of an electric field for which $\mathbf{J} \neq \mathbf{0}$, it is not clear how to represent such contributions. Finally, the specification of \mathbf{B} itself requires the presence of ghost nodes on the other side of the boundary, a clumsy admission of the method’s inadequacy near the boundary. Fortunately, this requirement is in line with that of our SPMHD algorithm.

4.3 A Meshless Petrov-Galerkin Approach to Magnetic Diffusion

As an example of how other mesh-free methods can be used to solve the magnetic diffusion equation, we now describe a approach based on the Meshless Local Petrov-Galerkin (MLPG) method [5] developed by Atluri et al as an alternative to the Element-Free Galerkin (EFG) method [50]. Galerkin-like methods are attractive for diffusion problems with complicated boundary conditions because these boundary conditions can often be expressed in terms of integral forms, avoiding the need for ghost elements or nodes. The EFG method is a variant of the familiar mesh-based Galerkin method that uses mesh-free shape functions that are based on distributions of nodes rather than finite elements, an idea used in the earlier Diffuse Elements method [48]. However, as Atluri observes, the EFG method is not truly mesh-free: the Galerkin integral forms are evaluated using quadrature on a grid. The MLPG method uses mesh-free quadrature rules and does not depend on a grid at all, making it a good candidate for combining with other mesh-free endeavors.

The main difference between the EFG method and the MLPG method is that the former is based on *global* weak forms, whereas the latter uses *local* ones. A global weak form is integrated over the entire domain of a problem, which allows one to take advantage of any special geometric features or symmetries of the domain but requiring a set of elements that completely cover the domain but do not overlap. Thus each

node contributes to the global form as a term in a sum. By contrast, a local weak form is an integral over a subdomain defined for each node. The subdomains of local weak forms are allowed to overlap, making it a natural fit for mesh-free methods, for which the resolution scale often spans several nodes. Below are examples of global and local weak forms and their corresponding domains.

We described the use of this method to treat magnetic diffusion in [52].

4.3.1 The Local Weak Form

Consider a conducting medium within a domain Ω . The medium has a dynamic symmetric-tensor-valued electrical resistivity $\eta(\mathbf{x}, t)$ that is related to temperature, density, and other properties. Within the medium is a magnetic field \mathbf{B} given at time $t = 0$ by $\mathbf{B}_0(\mathbf{x}, t)$. On the boundary $\Gamma = \Gamma_B \cup \Gamma_K \cup \Gamma_E$, the following boundary conditions hold:

$$\mathbf{B} = \mathbf{B}_\Gamma(\mathbf{x}, t) \text{ on } \Gamma_B$$

$$\mathbf{n} \times \mathbf{B} = \mathbf{K}(\mathbf{x}, t) \text{ on } \Gamma_K$$

$$\frac{1}{\mu_0} \mathbf{n} \times \eta \nabla \times \mathbf{B} = \mathbf{n} \times \mathbf{E}_\Gamma \text{ on } \Gamma_E$$

To devise a MLPG method for this problem, we introduce a set of computational points, or *nodes* within Ω at positions \mathbf{x}_i . A weak form for (4.19) can then be constructed over a set of local subdomains Ω_i corresponding to these points. The subdomain Ω_i represents the support of the i th node, and so the subdomains will typically overlap one another and cover the entire global domain Ω . In general, the subdomains may be of any shape or size; presently, we discuss circular subdomains for which the radius of the support of the i th point is given by R_i .

We write the desired *local* weak form by multiplying both sides of (4.19) by a test function $\psi_i(\mathbf{x})$ and integrating over Ω_i :

$$\int_{\Omega_i} \frac{d\mathbf{B}}{dt} \psi_i d\Omega = -\frac{1}{\mu_0} \int_{\Omega_i} \nabla \times (\boldsymbol{\eta} \cdot \nabla \times \mathbf{B}) \psi_i d\Omega \quad (4.40)$$

ψ can be any weight function or mesh-free shape function that has compact support on the subdomain Ω_i .

The right hand side of the equation can be integrated by parts using Green's Theorem for vectors to obtain

$$\int_{\Omega_i} \frac{d\mathbf{B}}{dt} \psi_i d\Omega = -\frac{1}{\mu_0} \int_{\Gamma_i} \mathbf{n} \times \boldsymbol{\eta} \cdot \nabla \times \mathbf{B} \psi_i d\Gamma + \frac{1}{\mu_0} \int_{\Omega_i} \boldsymbol{\eta} \cdot \nabla \times \mathbf{B} \cdot \nabla \times \psi_i d\Omega \quad (4.41)$$

Here, Γ_i is the intersection of the boundary Γ with the subdomain Ω_i for the i th node. In addition, we have defined a curl of the test function $\psi_i(\mathbf{x})$, which we have previously introduced as a scalar function. The expression for the “curl” of the test function is discussed in Appendix A.

The first term in the right hand side of (4.41) represents the contribution of a natural boundary condition on \mathbf{B} . If Γ_i intersects Γ_E , this term may be written as

$$-\int_{\Gamma_i} \mathbf{n} \times \mathbf{E}_\Gamma \psi_i d\Gamma$$

Otherwise, we leave the term in its original form and proceed with our analysis.

4.3.2 Semi-Discrete Equations

We wish to represent the vector-valued solution for the magnetic field on our set of nodes in Ω . The solution is approximated using a trial function $\phi(\mathbf{x})$ in the usual manner:

$$\mathbf{B}(\mathbf{x}, t) \approx \sum_{j \in \Omega_i} \phi_j(\mathbf{x}) \hat{\mathbf{B}}_j \quad (4.42)$$

where $\{\hat{\mathbf{B}}_j\}$ is a set of “fictitious” nodal values representing the projection of $\mathbf{B}(\mathbf{x}, t)$ onto the basis $\{\phi_j\}$.

ϕ can be any mesh-free shape function and is generally different from ψ (since this is a Petrov-Galerkin method). Substituting (4.42) into (4.41) and pulling the sums outside the integrals yields

$$\sum_j \int_{\Omega_i} \psi_i \phi_j \frac{d\hat{\mathbf{B}}_j}{dt} d\Omega = -\frac{1}{\mu_0} \sum_j \int_{\Gamma_i} \mathbf{n} \times (\boldsymbol{\eta} \cdot \nabla \times \phi_j) \psi_i \hat{\mathbf{B}}_j d\Gamma + \frac{1}{\mu_0} \sum_j \int_{\Omega_i} \boldsymbol{\eta} \cdot \nabla \times \phi_j \cdot \hat{\mathbf{B}}_j \cdot \nabla \times \psi_i d\Omega$$

Using the fact that $\boldsymbol{\eta}$ is either a scalar or a symmetric tensor, we observe that $\boldsymbol{\eta} \cdot \nabla \times \phi_j \cdot \hat{\mathbf{B}}_j = \hat{\mathbf{B}}_j \cdot \nabla \times \phi_j \cdot \boldsymbol{\eta}$ and so the last term in (4.43) becomes

$$\frac{1}{\mu_0} \sum_j \int_{\Omega_i} \hat{\mathbf{B}}_j \cdot (\nabla \times \phi_j \cdot \boldsymbol{\eta} \cdot \nabla \times \psi_i) d\Omega \quad \text{or} \quad \frac{1}{\mu_0} \sum_j \int_{\Omega_i} (\nabla \times \psi_i \cdot \boldsymbol{\eta} \cdot \nabla \times \phi_j) \cdot \hat{\mathbf{B}}_j d\Omega \quad (4.43)$$

(4.43) is the MLPG semi-discrete form of the magnetic diffusion equation on the subdomain Ω_i . It is a vector valued equation with 3 components: one for each of the components of the magnetic induction \mathbf{B} . At this point, we can follow the common practice for Galerkin and Petrov-Galerkin methods and express this vector equation as a linear system, interpreting $\{\hat{\mathbf{B}}_j\}$ as the components of a solution vector \mathcal{B} :

$$\mathbf{M} \frac{d\mathcal{B}}{dt} = \frac{1}{\mu_0} \mathbf{K}(\boldsymbol{\eta}) \mathcal{B} + \frac{1}{\mu_0} \mathbf{f} \quad (4.44)$$

where \mathbf{M} , the “mass matrix,” \mathbf{K} , the “stiffness matrix,” and \mathbf{f} , the “boundary vector” are defined as follows:

$$\begin{aligned} \mathbf{M}_{ij} &\equiv \int_{\Omega_i} \psi_i \phi_j d\Omega \\ \mathbf{K}(\boldsymbol{\eta})_{ij} &\equiv \int_{\Omega_i} \nabla \times \psi_i \cdot \boldsymbol{\eta} \cdot \nabla \times \phi_j d\Omega - \int_{\Gamma_i \cup (\Gamma_B \cup \Gamma_K)} \psi_i (\mathbf{n} \times \nabla \times \phi_j) d\Gamma \\ \mathbf{f} &\equiv - \int_{\Gamma_i \cup \Gamma_E} \mathbf{n} \times \mathbf{E}_\Gamma \psi_i d\Gamma \end{aligned} \quad (4.45)$$

The matrices \mathbf{M} and \mathbf{K} are 3×3 matrices, and \mathcal{B} and \mathbf{f} are 3-vectors. If the domain Ω contains N nodes, the global linear system will thus comprise N linear systems described by (4.44) and the solution vector \mathbf{b} will have $3N$ values as expected. Note that we have also used the “curl” of the trial function ϕ . Its evaluation is analogous to that of $\nabla \times \phi$.

As expected for a MLPG method, \mathbf{M} and \mathbf{K} are not symmetric matrices ([133], [5]). Consequently, we do not expect that the fully discrete global linear system will have a symmetric matrix. However, we have already seen that this is no great obstacle for the SPH magnetic diffusion method. Specific forms of the matrices are given in Appendix A.

4.3.3 Meshless Shape Functions

We have mentioned that the trial function ϕ and the test function ψ can take various forms. One of the great strengths of the MLPG method is that it allows one to choose these functions from a larger cast of mesh-free shape functions and weight functions that have compact support. These typically include Shepard functions, Moving Least Squares (MLS) functions, Radial Basis Functions (RBF), and others. These shape functions are extremely popular not only in MLPG methods [133] but also in several other mesh-free methods ([50], [85]). There are several excellent expositions of these shape functions in the above references and also in [51] and [134]; for completeness, we briefly describe the MLS shape functions, which we employ in the present method.

Consider a domain Ω in space in which a continuous function F is defined at several discrete points $\{\mathbf{x}_i\}$. To approximate the value of F at some point \mathbf{x} within this domain, we define the MLS approximant (itself a continuous function of space) as

$$\tilde{F}(\mathbf{x}) = \mathbf{p}^T(\mathbf{x})\mathbf{a}(\mathbf{x}), \quad \mathbf{x} \in \Omega \quad (4.46)$$

where $\mathbf{p}^T(\mathbf{x})$ is some polynomial basis of order m (i.e. $[p_1(\mathbf{x}), p_2(\mathbf{x}), \dots, p_m(\mathbf{x})]$) and $\mathbf{a}(\mathbf{x})$ is the corresponding

projection of F onto that basis. Some examples of $\mathbf{p}^T(\mathbf{x})$ for various dimensions and orders appear below:

$$\mathbf{p}^T(\mathbf{x}) = [1, x], \text{ 1D, linear } (m = 2)$$

$$\mathbf{p}^T(\mathbf{x}) = [1, x, x^2], \text{ 1D, quadratic } (m = 3)$$

$$\mathbf{p}^T(\mathbf{x}) = [1, x, y], \text{ 2D, linear } (m = 3)$$

$$\mathbf{p}^T(\mathbf{x}) = [1, x, y, x^2, xy, y^2], \text{ 2D, quadratic } (m = 6)$$

$$\mathbf{p}^T(\mathbf{x}) = [1, x, y, z], \text{ 3D, linear } (m = 4)$$

$$\mathbf{p}^T(\mathbf{x}) = [1, x, y, z, x^2, xy, y^2, yz, z^2, xz], \text{ 3D, quadratic } (m = 10)$$

To project F onto the desired basis, one assumes a set of fictitious nodal values $\{\hat{F}_i\}$ that exist on all points $\{\mathbf{x}_i\}$ within Ω and then proceeds to minimize a weighted L_2 norm defined by

$$J(\mathbf{x}) = \sum_{i \in \Omega} W_i(\mathbf{x}) [\mathbf{p}^T(\mathbf{x}_i) \mathbf{a}(\mathbf{x}) - \hat{F}_i]^2 \quad (4.47)$$

where $W_i(\mathbf{x})$ is some compactly-supported weight function centered about the point \mathbf{x}_i in the domain Ω . In general, $\hat{F}_j \neq \tilde{F}(\mathbf{x}_j)$. Rather, the relationship between the fictitious and the “actual” nodal values can be expressed in terms of the desired shape functions $\{\phi_i(\mathbf{x})\}$ as in (A.1).

Supposing that there are N points within Ω and that the dimension of our polynomial basis is m , the problem of minimizing (4.47) (and thus projecting F onto the basis \mathbf{p}^T) is equivalent to solving a least-squares problem described by the $m \times m$ linear system

$$\mathbf{A}(\mathbf{x}) \mathbf{a}(\mathbf{x}) = \mathbf{B}(\mathbf{x}) \hat{\mathbf{F}} \quad (4.48)$$

where the matrices \mathbf{A} and \mathbf{B} are defined by

$$\mathbf{A}(\mathbf{x}) = \sum_{i=1}^N W_i(\mathbf{x}) \mathbf{p}(\mathbf{x}_i) \mathbf{p}^T(\mathbf{x}_i) \quad (4.49)$$

$$\mathbf{B}(\mathbf{x}) = [W_1(\mathbf{x}) \mathbf{p}(\mathbf{x}_1), \dots, W_N(\mathbf{x}) \mathbf{p}(\mathbf{x}_N)] \quad (4.50)$$

and $\hat{\mathbf{F}}$ is the vector of fictitious nodal values $\{\hat{F}_i\}$. Note that \mathbf{A} and \mathbf{B} are $m \times m$ and $m \times N$ matrices, respectively. Thus, $\mathbf{a}(\mathbf{x})$, the projection of F onto $\mathbf{p}^T(\mathbf{x})$, becomes $\mathbf{A}^{-1}(\mathbf{x})\mathbf{B}(\mathbf{x})\hat{\mathbf{F}}$ and we obtain the MLS shape functions $\{\phi_i(\mathbf{x})\}$, defined by

$$\phi_i(\mathbf{x}) = \sum_{j \in \Omega} p_j(\mathbf{x})[\mathbf{A}^{-1}(\mathbf{x})\mathbf{B}(\mathbf{x})]_{ji} \quad (4.51)$$

The *moment matrix* \mathbf{A} is symmetric, and so as long as it is nonsingular, (4.48) can be solved using Cholesky factorization. In practice, one can ensure that \mathbf{A} is positive definite by choosing sufficient numbers of points in the neighborhood of \mathbf{x} . A rule of thumb adopted by many practitioners is to choose m as roughly twice the number of points as is needed to find a unique solution. This helps to ease difficulties encountered in regions where F is not smooth.

The partial derivatives of these shape functions can be obtained directly using the expression

$$\frac{\partial \phi_i}{\partial x_k} = \sum_{j=1}^q \left[\frac{\partial p_j}{\partial x_k} (\mathbf{A}^{-1} \mathbf{B})_{ji} + p_j (\mathbf{A}^{-1} \frac{\partial \mathbf{B}}{\partial x_k} + \frac{\partial \mathbf{A}^{-1}}{\partial x_k} \mathbf{B})_{ji} \right] \quad (4.52)$$

(where $\frac{\partial \mathbf{A}^{-1}}{\partial x_k} = -\mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial x_k} \mathbf{A}^{-1}$) [133], or by using the “diffuse derivatives” explored by [48]. Presently, we directly compute the derivatives using (4.52).

As we have mentioned, each mesh-free shape function is constructed using a weight function with compact support [133]. Our implementation uses a quartic spline function identical to that used in several previous studies of MLPG methods:

$$W_i(\xi) = \begin{cases} 1 - 6\xi^2 + 8\xi^3 - 3\xi^4 & 0 \leq \xi < 1, \\ 0 & \xi \geq 1. \end{cases} \quad (4.53)$$

Here, $\xi \equiv |\mathbf{x} - \mathbf{x}_i|/R_i$, where R_i is the radius of support of W_i about \mathbf{x}_i .

Atluri et al [5] have found that the convergence rates for Shepard functions and MLS functions exceed h -convergence, which is given for computed and exact solutions u_c and u_e :

$$\|u_c - u_e\| \propto h^{m+1-k} \quad (4.54)$$

where $\|\cdot\|_k$ is the k -norm in the Sobolev space and m is the order of completeness in the polynomial interpolation. We have verified that h -convergence is indeed a conservative estimate for the convergence of the interpolation error of the mesh-free shape functions. The MLPG method, like any weak form, is based upon the integration of interpolated quantities, so the convergence rates of the meshless interpolation functions represent a characteristic rate of convergence in the spatial discretization for methods that use them.

The MLPG approach is, in the most general sense, a family of related methods. The members of this family are defined by the choice of the test function ψ , which is allowed to differ in form from the trial function ϕ . Atluri classifies six variants of MLPG, labeling them numerically as MLPG1-6 [51]. Each one of these variants prescribes a relationship between these two functions. We focus our attention primarily on the first and fifth variants. MLPG1, the first, specifies that ϕ is one of the aforementioned mesh-free shape functions, and that ψ is the aforementioned weight function W that was used in the construction of ϕ . This is the original version of the MLPG method and produces accurate results, but requires many quadrature points to evaluate the integrals because of the complicated structure of the integrands. The other variant, MLPG5, uses the Heaviside function as the test function, producing greatly simplified integrals that require fewer quadrature points for accuracy. Another variant, MLPG2, avoids numerical integration altogether by using a Dirac Delta function as the test function. This is the variant that is closest in spirit to SPH, but suffers from complexities in the second derivatives of the meshless shape functions, since no integration by parts is performed in the derivation of the weak form.

4.3.4 Fully Discrete Equations

The system (4.44) can be made totally discrete in a manner very similar to that of (4.28), resulting in the linear system

$$\left[\mathbf{M} - \theta \Delta t \mu_0^{-1} \mathbf{K}(\boldsymbol{\eta}^{n+\theta}) \right] \mathcal{B}^{n+1} = \left[\mathbf{M} + (1 - \theta) \Delta t \mu_0^{-1} \mathbf{K}(\boldsymbol{\eta}^{n+\theta}) \right] \mathcal{B}^n \quad (4.55)$$

using the same implicitness parameter θ . The same conditions that applied to the SPH discretization apply to this system. Like the SPH method, the structure of the matrix in (4.55) is unsymmetric, and so we use the

same GMRES solver and Jacobi preconditioner to obtain our solutions.

4.3.5 Boundary Conditions

In this method, the boundary conditions are treated more explicitly: they appear directly in the weak forms (4.41) and (4.43). The term referring to \mathbf{E} covers the specification of Neumann boundaries.

Dirichlet Boundaries

Since our method is based on local weak forms, the values of the solution incorporate all points within their subdomains. This makes it difficult to enforce essential boundary conditions directly even when trial functions possessing the Dirac delta function property are used [134]. One way to indirectly impose values on the boundary Γ of the domain Ω is to use a penalty method with some multiplier $\alpha \gg 1$. This adds the effect of a strong spring on the boundary that holds the solution at the prescribed value. In this case, penalty terms

$$\alpha \int_{\Gamma_i} \psi_i(\mathbf{x}) \phi_j(\mathbf{x}) d\Gamma \quad \text{and} \quad \alpha \int_{\Gamma_i} \mathcal{B}_\Gamma(\mathbf{x}) \psi_i(\mathbf{x}) d\Gamma \quad (4.56)$$

are added to the left and the right hand sides of (4.43), respectively. A problem with this approach is that α must be “big enough” so that the penalty terms overshadow the others, but not so large that it renders the matrix ill-conditioned. In other words, α is a free parameter that can be sensitive to the nodal spacing in the problem, meaning that it must vary with the resolution. Nevertheless, it is a straightforward approach that gives convergent solutions. We have found that values of α between 10^6 and 10^8 give the best results when using MKS units (in which $\mu_0 = 4\pi \times 10^{-7} \text{N} \cdot \text{A}^2$).

Other approaches have been suggested. The *transformation* method, described in [5], imparts the Dirac delta function property to the solution at the boundary via a transformation of the linear system (4.55). The transformation matrix $\{R_{ij}\} = \{\phi_j(\mathbf{x}_i)\}^{-1}$ can be computed using LU decomposition to invert the matrix of trial functions, but care must be taken to mitigate numerical errors and the expense of transforming the

system. A *modified colocation* method, described in [135], uses the interpolation

$$\mathbf{B}(\mathbf{x}_i, t) = \sum_{j \in \Omega_i} \phi_j(\mathbf{x}_i) \hat{\mathbf{B}}(t)$$

to replace each row of (4.55) corresponding to a node on Γ , with the right hand side set to the value prescribed by the function \mathbf{B}_Γ .

$$\sum_{j \in \Omega_i} \phi_j(\mathbf{x}_i) \hat{\mathbf{B}}(t) = \mathbf{B}_\Gamma(\mathbf{x}_i, t), i \in \Gamma \quad (4.57)$$

Here, one must be careful to eliminate all other entries on these boundary rows in order to avoid introducing spurious errors on the boundary. More recently, [136] have described the application of this method to their “mixed” finite-volume MLPG approach in elasticity problems.

We have used both the penalty method and the modified colocation method and have found that both give convergent results for our test problems. However, the treatment of essential boundary conditions for this method remains an open issue, as surfaces in two and three dimensions are invariably more complicated. For both Neumann and Dirichlet boundaries that are specified as analytic functions, one must evaluate these functions at the properly-centered times to achieve the expected rates of convergence.

Periodic Boundaries

For periodic and asymmuthally-symmetric boundary conditions, we use the same node-aliasing strategy that we did for the SPH system. The structure of the matrix in the linear system is topologically identical to that of the SPH system—the (i, j) th component of the matrix refers to an interaction between the same two nodes in each system—so this strategy will work as is.

4.3.6 Limitations

The power and flexibility of the MLPG method is evident (and we demonstrate its solving power in the test problems where relevant) but comes at a price: the method is extremely computationally expensive when the nodes are allowed to move. This is largely because moving the nodes (and changing the topology)

requires not only a neighbor search, but the recalculation of the matrices, and therefore of the mesh-free shape functions at each of the quadrature points in the domain. In fact, the cost of the neighbor search is all but lost in the noise of the quadrature computation in 2D and 3D.

We should reiterate at this point that the connectivity in the MLPG method is much more rich than that of SPH. Each quadrature point sees a set of nodes whose shape functions are evaluated at that point. If there are n quadrature points per node and m nodes, and each quadrature point sees p neighboring nodes, then the number of mesh-free function evaluations is mnp ! If we use a stencil “equivalent” to an SPH stencil with n_h , then $p \approx 50$ in 2D. With a modest 10-point quadrature rule on a 100×100 lattice of nodes, we have to evaluate $100 \times 100 \times 10 \times 50 = 5,000,000$ shape function values every time we update a matrix!

Interestingly, Atluri makes some simple estimates of performance that place the method into a competitive position with the finite element method [5]. His argument is that the number of nodes for the MLPG to achieve a desired level of accuracy corresponds roughly to an expense at which the same accuracy is achieved with finite elements. Of course, finite element methods also suffer from their expense when mass and stiffness matrices are repeatedly computed, for essentially the same reasons as does the MLPG method.

We have gone to some length to try to reduce the cost of the method by caching quadrature points, shape function values, and topology, but have been unable to make the method competitive with SPH. Thus it is not particularly surprising that, in the mesh-free community, SPH and similar “strong-form” methods have retained their dominance in Lagrangian fluid calculations. However, given its accuracy in expressing Neumann boundary conditions on magnetic fields, the MLPG method we have described holds promise in the evaluation of vacuum magnetic fields on fixed sets of points. We will explore this idea in a future publication.

4.4 Resistive Smoothed Particle Magnetohydrodynamics

We have demonstrated that we know how to solve the magnetic diffusion problem. It remains to combine this technique with SPMHD, which we discussed in the last chapter. Fortunately, these two methods have a lot in common, and this combination is fairly simple. First we compare the stencils of the two methods

to see how well they might run together. We then discuss an easy way to apply the different contributions to **B**.

Ironically, adding magnetic diffusion to the MHD equations can make things easier. The equations are numerically more stable, since no artificial dissipation is needed to smooth **B**. The decay of the magnetic field itself also eases some of the effects of the tensile instability. And the conservation of magnetic flux need no longer be rigidly obeyed. We explore the effects of inserting magnetic diffusion into ideal MHD test problems at the end of the chapter.

4.4.1 Stencil Coupling

Both of the mesh-free magnetic diffusion methods we have described have parameters defining their characteristic “stencils.” In SPH, the radius of visibility of a node i is its smoothing length h_i times the extent κ of the smoothing kernel (which is 2 for the cubic B-spline). Recall that the smoothing length is related to the particle spacing Δx_i by the number of nodes per smoothing scale, n_h .

On the other hand, the radius of each MLPG trial function is specified in terms of a number of particle spacings ρ (not to be confused with the mass density), so that $r_i^{Tr} = \rho \Delta x_i$. This suggests that the parameters ρ and n_h are linked. If we equate the trial function radius to the radius of visibility of an SPH node, we obtain the relation

$$\rho = \kappa n_h \tag{4.58}$$

(or simply $\rho = 2n_h$ for the B-spline kernel). The radius of an MLPG test function is related to that of the trial function by the ratio r^{Te}/r^{Tr} , and if we wish to preserve the practice that $r_i^{Te} = \Delta x_i$, we can set r^{Te}/r^{Tr} to $1/\rho = 1/2n_h$.

Happily, our preferred range of n_h values between 1.5 and 2 correspond to values of ρ between 3 and 4, a range in which MLPG has been demonstrated to perform very reliably.

4.4.2 Operator Splitting

We originally discussed the magnetic induction equation in its most general form in Chapter 2. If we incorporate anisotropic resistance, this equation becomes

$$\frac{D\mathbf{B}}{Dt} = \mathbf{B} \cdot \nabla \mathbf{v} - \mathbf{B} (\nabla \cdot \mathbf{v}) + \nabla \times \left(\frac{\mathbf{R}}{\mu_0} \cdot \nabla \times \mathbf{B} \right) \quad (4.59)$$

We are already capable of evolving the two terms in $\frac{D\mathbf{B}}{Dt}$ separately. The most straightforward way to combine these terms into one is to assume that neither process affects the other over a time interval Δt so that, for node i ,

$$\frac{D\mathbf{B}_i}{Dt} = \left(\frac{D\mathbf{B}_i}{Dt} \right)_{\text{MHD}} + \left(\frac{\partial \mathbf{B}_i}{\partial t} \right)_{\text{diffusion}} \quad (4.60)$$

We compute the first term according to (4.3). The second one is computed by a simple finite difference in time after solving the magnetic diffusion equation:

$$\left(\frac{\partial \mathbf{B}_i}{\partial t} \right)_{\text{diffusion}} = \frac{\mathbf{B}_i^{n+1} - \mathbf{B}_i^n}{\Delta t} \quad (4.61)$$

In this way, any time integration scheme can be used to update \mathbf{B} at every node consistently. If a multi-stage time integrator is used, the magnetic diffusion equation should be solved at each integration stage for the best accuracy. It may be possible to avoid this expense by computing $\left(\frac{\partial \mathbf{B}}{\partial t} \right)_{\text{diffusion}}$ only once per time step, but this can result in a larger operator splitting error. Since we have shown that our single-stage second-order Leapfrog integrator obtains convergent results, we can sidestep this potential complication.

4.5 Test Problems

We conclude this chapter with selected test problems to demonstrate the capabilities of the techniques. These problems cover much ground, from the propagation of different MHD waves to turbulent vortices to the diffusion of magnetic flux tubes. We have performed formal convergence studies wherever

possible. It is unfortunate that this appears to be the first time that convergence rates have been published for any SPH method treating magnetic fields.

Unfortunately there are relatively few problems in MHD that have known solutions. Because of this, we have verified many of our results against those given by the Godunov code *Athena*, complements of Jim Stone [40]. As we mentioned in Chapter 3, it is common practice to compare one’s answers to a “highly-resolved” Riemann-based calculation in the absence of analytic solutions. The demonstration of “self-convergence” on the part of the Riemann solver (that it converges to *an* answer) seems sufficient for the definition of such a benchmark. The fishiness of this methodology speaks more to the difficulty of finding good MHD test problems than to anything else.

The test problems for the MLPG method appear in [52] and are not covered in this chapter, since these results do not ultimately factor into the simulations in Chapter 5.

The simplest problems with magnetic fields involve the decay of such fields within imperfect conductors. In these problems there is no motion, and the dynamics is much simpler than that of MHD. The nodal mass and mass density is immaterial in these problems, so we use $\rho = 1$ with the corresponding nodal mass $m_0 = \rho V/N = V/N$, where V is the volume of the system and N is the total number of nodes. In problems without discontinuities, we expect the error in our solution to converge at second-order in Δx and Δt when the Crank-Nicholson time integrator is used.

4.5.1 The decay of a magnetic pulse

Jackson [62] describes the decay of a square magnetic pulse within an infinite conducting medium with magnetic diffusivity D . Consider a conducting medium on the domain $[-L/2, L/2]$. If B_y is initially 1 on $[-L/4, L/4]$ and 0 elsewhere, its value at a subsequent time t is given by

$$B_y(x, t) = \operatorname{erf}\left(\frac{2(1+|x|)}{4\sqrt{Dt}}\right) + \operatorname{erf}\left(\frac{2(1-|x|)}{4\sqrt{Dt}}\right) \quad (4.62)$$

where erf denotes the error function

$$\operatorname{erf}(z) \equiv \frac{2}{\sqrt{\pi}} \int_0^z e^{-u^2} du \quad (4.63)$$




Figure 4.1: Initial configuration for (a) \mathbf{B} and (b) \mathbf{J} in the decay of a square magnetic pulse.

Thus the pulse begins as the square shown in Figure 4.1 and spreads over time, assuming a bell shape. Since the initial profile is discontinuous, we expect only first-order convergence. We run this problem with numbers of nodes $N = 100$, $N = 200$, $N = 400$, and $N = 800$ on a domain $[0, 1]$ with $D = 0.01$. We use the backward-Euler integrator ($\theta = 1$) to guarantee that the solution is stable in the presence of discontinuities, and we choose a constant time step $\Delta t = 0.01\Delta x/D$ for each resolution. We use ghost nodes on either side of the domain to enforce Dirichlet boundary conditions on \mathbf{B} . Since the nodes are stationary in this and other pure diffusion problems, the tensile instability poses no threat and we use $n_h = 2$ to compute these results. The computed solutions for \mathbf{B} and \mathbf{J} at time $t = 0.5$ are shown in Figure 4.2 as dots superimposed on the solid lines representing the analytic profiles. The convergence of these solutions is shown in Figure 4.3. We see that the expected first-order convergence is actually exceeded in the solution for \mathbf{B} , and that the accuracy is very good even for relatively low resolutions.

If the initial profile is instead a Gaussian pulse, there are no discontinuities and we expect to achieve second-order convergence. The decay of a Gaussian pulse is simply a wider Gaussian, and the analytic




Figure 4.2: Computed and analytic solutions for (a) \mathbf{B} and (b) \mathbf{J} in the decay of a square magnetic pulse at $t = 0.5$.




Figure 4.3: Convergence of solutions for (a) \mathbf{B} and (b) \mathbf{J} in the decay of a square magnetic pulse at $t = 0.5$.

solution as a function of t is

$$B_y(x, t) = \frac{1}{\sqrt{4\pi a D t}} \exp\left(\frac{-x^2}{4a}\right) \quad (4.64)$$

where a is a number (e.g. 10^{-3}) characterizing the width of the pulse. We run this problem on the same domain with $D = 0.01$, showing the solutions at $t = 0.5$ in Figure 4.4 and their convergence in Figure 4.5.

Using the Crank-Nicholson time integrator and the same time step we used for the square pulse, we achieve the desired second-order rate of convergence.




Figure 4.4: Computed and analytic solutions for (a) \mathbf{B} and (b) \mathbf{J} in the decay of a Gaussian magnetic pulse at $t = 0.5$.

To get an idea of how the technique extends to multi-dimensional problems, we run a simulation of a square pulse in the $x - y$ plane diffusing outward. This is analogous to our 1D problem, and we show the magnitude of \mathbf{B} at $t = 0$ and at $t = 0.5$ in Figure 4.6. The spread of the pulse is equivalent to that in the 1D problem. We approach 2D magnetic diffusion more rigorously in an another test problem.

Unlike the hydrodynamics algorithm, this magnetic diffusion algorithm appears not to be sensitive to n_h , and achieves the same convergence rates for \mathbf{B} and \mathbf{J} over a large range of values for this parameter. Therefore, we will not pay close attention to n_h for pure magnetic diffusion problems from this point.




Figure 4.5: Convergence of solutions for (a) \mathbf{B} and (b) \mathbf{J} in the decay of a Gaussian magnetic pulse at $t = 0.5$.




Figure 4.6: (a) Initial ($t = 0$) and (b) final ($t = 0.5$) solutions for $|\mathbf{B}|$ in the two-dimensional decay of a square magnetic pulse.

4.5.2 The magnetic “wine cellar” problem

As a test of our formulation of time-dependent Dirichlet boundary conditions, we model the penetration of an oscillating magnetic field into a real conductor. This problem is the magnetic analog to a well-known problem in which the heat equation is used to decide how far underground it is best to build a wine cellar in order to quell the temperature variation induced by the rising and setting sun at the Earth’s surface. If we impose an external field such that $\mathbf{B}(0,t) = (0,0,B_0 \sin \omega t)$, the analytic solution on a 1D domain $[0,L]$ is

$$B_z(x,t) = \Re \left\{ \frac{\sin(\beta(L-x))}{\sin(\beta L)} \exp(-i\omega t) \right\}$$

The more resistive the material, the further the fluctuations in the field will penetrate. We use ghost nodes to impose the external field at $x = 0$ and the field $\mathbf{B} = \mathbf{0}$ at $x = L = 1$, and initialize the magnetic field to the solution at time 0 with $\omega = 1$. Using the resolutions $N = 100, N = 200, N = 400, N = 800$ and a diffusivity $D = 0.01$ with the Crank-Nicholson time integrator and a time step size $\Delta t = 0.001\Delta x/D$, we obtain the solutions for \mathbf{B} and \mathbf{J} shown in Figure 4.7 at time $t = 0.1$. Figure 4.8 shows the convergence of these solutions. We achieve second-order convergence for this problem.

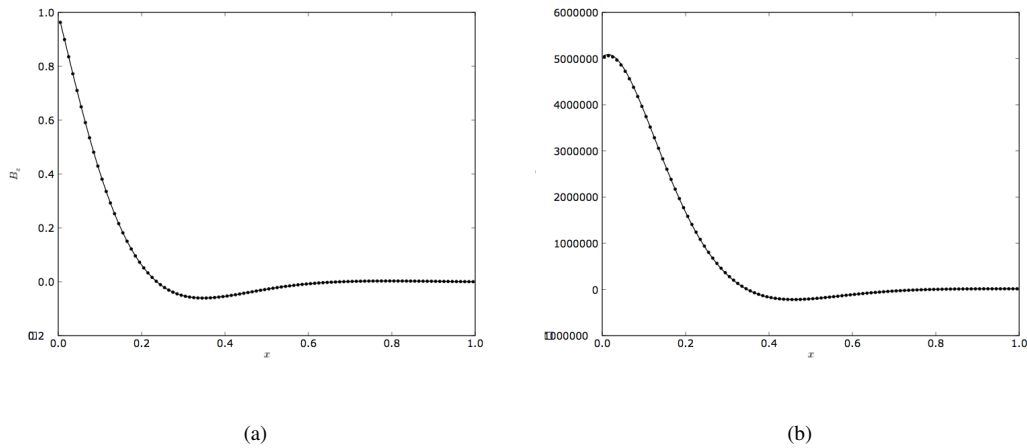


Figure 4.7: Computed and analytic solutions for (a) \mathbf{B} and (b) \mathbf{J} in the magnetic “wine cellar” problem at $t = 0.1$.




Figure 4.8: Convergence of solutions for (a) \mathbf{B} and (b) \mathbf{J} in the magnetic “wine cellar” problem at $t = 0.1$.

We have mentioned the insensitivity to nodal disorder as a distinct advantage to the Brookshaw SPH Laplacian operator, and have performed a convergence study on the solutions for \mathbf{B} and \mathbf{J} to this problem for nodes that are randomly shuffled by up to 20% of Δx . The results are shown in Figure 4.9. The convergence is degraded, but is still well above first order. Finite-element-like methods that use the weak form of the partial differential equations (such as the MLPG method) tend to be less sensitive to point disorder than their strong-form counterparts, so these results speak well for the treatment of magnetic diffusion in messy magnetohydrodynamic problems.

4.5.3 Continuously-varying resistivity

If we are to model plasmas with realistic resistivity models, it is important that we establish that the technique can handle resistivities that are not uniform in space. In an unpublished paper by Wilmot-Smith [137], the analytic solution for a magnetic field within a conductor with a continuously-varying diffusivity. The diffusivity has the quadratic form

$$D(x) = \frac{D_0}{L^2}(L^2 - x^2) \quad (4.65)$$




Figure 4.9: Convergence of solutions for (a) \mathbf{B} and (b) \mathbf{J} in the magnetic “wine cellar” problem with perturbed nodes at $t = 0.1$.

and is defined on a domain $[0, L]$. However, we must be careful not to include the region around $x = L$ in our simulation, since the diffusivity goes negative at that point. For safety, we run this problem on the domain $[0, L/2]$. An analytic solution to the magnetic diffusion equation in this case is

$$B_y(x, t) = B_0 \exp(-2D_0 t / L^2) (x/L). \quad (4.66)$$

Since D varies in space, we must be sure to assign the right values of D to the ghost nodes representing the domain boundary using (4.65). Otherwise erroneous discontinuities in D at the boundary will cause oscillations in the solution. We run this problem with $D_0 = B_0 = L = 1$ using resolutions $N = 100, N = 200, N = 400$, and $N = 800$. We use $\alpha = 0.5$ and choose a constant step size $\Delta t = 0.01 \Delta x / D_0$ to reduce the time discretization error. The computed values of \mathbf{B} at $t = 0.01$ and the convergence of the error are shown in Figure 4.10.

The method resolves \mathbf{B} very nicely in the presence of the varying resistivity, and the convergence in the error is essentially second-order. This demonstrates that the method is capable of handling spatially inhomogeneous resistivity, provided that it is continuous.

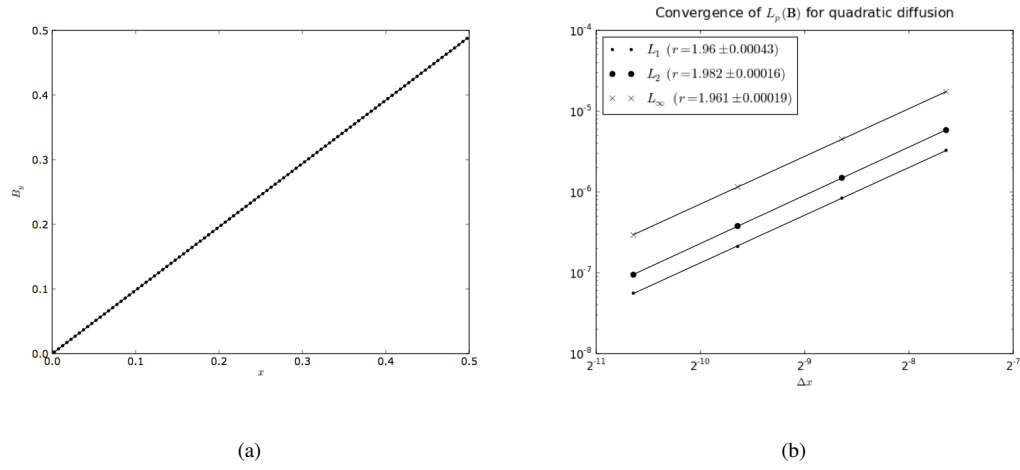


Figure 4.10: (a) Solution for \mathbf{B} in a conductor with continuously-varying diffusivity at $t = 0.1$ and (b) the convergence of its L_p error norms for $N = 100, N = 200, N = 400, N = 800$.

4.5.4 Discontinuous resistivity

We now examine a problem in which the diffusivity is piecewise-constant. Two conductors lay adjacent to one another, the first one having diffusivity D_1 and the second D_2 . If the first conductor fills a region $[0, L]$ and the second fills $[L, a]$, then the magnetic diffusion equation is satisfied by a magnetic field of the form

$$B_y(x, t) = B_0 \begin{cases} \frac{x}{a} \left[\frac{(D_2/D_1)(a/L)}{(D_2/D_1) + (a/L) - 1} \right] + \alpha \exp \left(-\Lambda^2 \frac{D_1 t}{a^2} \right) \sin \left(\Lambda \frac{x}{a} \right), & 0 < x < L, \\ \left(\frac{x}{a} - 1 \right) \left[\frac{(a/L)}{(D_2/D_1) + (a/L) - 1} \right] + 1 + \\ \alpha \exp \left(-\Lambda^2 \frac{D_1 t}{a^2} \right) \frac{\sin \left(\Lambda \frac{L}{a} \right)}{\sin \left[\Lambda \sqrt{\frac{D_1}{D_2}} \left(\frac{L}{a} - 1 \right) \right]} \sin \left[\Lambda \sqrt{\frac{D_1}{D_2}} \left(\frac{x}{a} - 1 \right) \right], & L < x < a. \end{cases} \quad (4.67)$$

where α is an arbitrary constant and Λ is any solution of

$$\tan \left(\Lambda \frac{L}{a} \right) = \sqrt{\frac{D_1}{D_2}} \tan \left[\Lambda \sqrt{\frac{D_1}{D_2}} \left(\frac{L}{a} - 1 \right) \right] \quad (4.68)$$

We run this problem in a similar way to the last two configurations, with $\alpha = 1$ to accommodate the discontinuity. We use $\Lambda = 4.373$, and piecewise constant magnetic diffusivities $D_1 = 1$ and $D_2 = 4$. Once

again, we use a constant step size $\Delta t = 0.01\Delta x/D_2$. The solution for \mathbf{B} at $t = 1$ and the convergence of its error are shown in Figure 4.11.




Figure 4.11: (a) Solution for \mathbf{B} in a region with discontinuous diffusivity at $t = 1$ using (4.26), and (b) the convergence of its L_p error norms for $N = 100, N = 200, N = 400, N = 800$.

The discontinuity in D occurs at $x = 0.5$, and it is here that the computed solution deviates the most from the answer. This error results from our assumption in (4.25) that D is differentiable everywhere. Evidently, the error in this region does not decrease with higher resolution. This indicates that the method does not handle adjacent regions with significantly different resistivities gracefully.

If we instead use the product rule with SPH gradients to approximate $\nabla \cdot (D\nabla\mathbf{B})$ with (3.23) for ∇D , we still see an error near the discontinuity. Here, though, the computed solution overshoots the actual answer rather than undershooting it as our first attempt did, and the error converges, albeit very slowly. These results are shown in Figure 4.12.

While this method is not satisfactory for treating discontinuities in D , it is arguably a small improvement. This method is not typically used in SPH simulations of scalar diffusion because it requires that ∇D be computed, which involves an extra traversal of the neighbors of each node. However, this incurs no extra cost to us, since we must already compute this gradient. Note that this method is not a general-purpose replacement. When D is continuous, it yields much less convergent results than the previous method, so it should be used




Figure 4.12: (a) Solution for \mathbf{B} in a region with discontinuous diffusivity at $t = 1$ using the product rule and (3.23) for $\nabla \cdot (D\nabla \mathbf{B})$, and (b) the convergence of its L_p error norms for $N = 100, N = 200, N = 400, N = 800$.

sparingly, if at all.

4.5.5 Time-dependent resistivity

We have also claimed that the method handles resistivity that changes in time. The next test problem uses a dynamic resistivity model that does not vary in space. The resistivity decays exponentially according to the prescription

$$D(t) = D_0 \exp(-t/\tau) \quad (4.69)$$

over a domain $[-L, L]$, where τ is a characteristic rate of decay. In this case an analytic solution is given by the sum

$$B_y(x, t) = \sum_n \frac{2B_0}{n\pi} \exp\left[\frac{-n^2\pi^2 D_0 \tau}{L^2} (1 - e^{t/\tau})\right] \sin\left[\frac{n\pi x}{L}\right]. \quad (4.70)$$

We compute the solution using $\alpha = 0.5$ and set $D_0 = L = \tau = 1$. As we did in the last two problems, we set $\Delta t = 0.01\Delta x/D_0$. The solution for \mathbf{B} at $t = 0.1$ and its convergence are shown in Figure 4.13. We achieve

the desired second-order convergence, indicating that resistivities that vary in time are handled readily by the method. In particular, this verifies that the centering of $D(t)$ is handled properly by our implicit integration method.

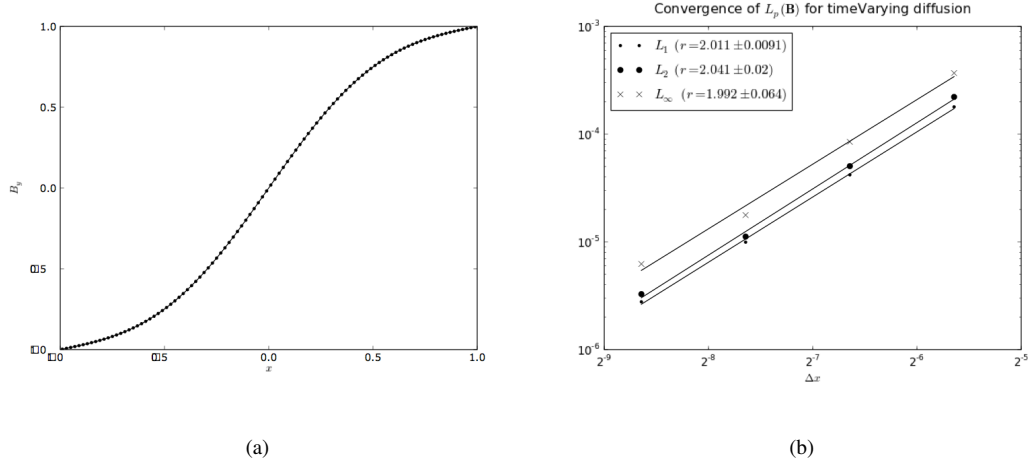


Figure 4.13: (a) Solution for \mathbf{B} for time-varying diffusivity $D(t) = D_0 e^{-t/\tau}$ at $t = 0.01$ and (b) the convergence of its L_p error norms for $N = 100, N = 200, N = 400, N = 800$.

4.5.6 Decay of a 2D magnetic flux tube

We now wish to test the diffusion of a magnetic field in two dimensions. An interesting 2D problem given by Wilmot-Smith [138] is the magnetic diffusion of a field within a cylindrical flux tube. In a cylindrical medium of diffusivity D , a self-similar solution to the magnetic diffusion equation is given by

$$B_\theta = \frac{r}{2Dt^2} \exp(-r^2/4Dt). \quad (4.71)$$

The other components of \mathbf{B} are zero. We begin the simulation at $t = 1$ since the solution is not finite at $t = 0$. We study this system at resolutions of $N = 40, N = 80, N = 160$, and $N = 320$ nodes on a side and starting from an initial time $t = 1$. We use a magnetic diffusivity of $D = 1$, and a constant time step size $\Delta t = 0.1\Delta x/D$. For expediency, we use the azimuthal boundary condition described in Chapter 3 to run the problem in the first quadrant. This means the nodes near the axis and near the quadrant boundaries will be replicated (and

rotated) in each quadrant. Recall that we modeled axisymmetric hydrodynamic systems in the last chapter by using simple rigid boundaries at $x = 0$ and $y = 0$. We cannot use this trick here, since that would produce the wrong values for B_θ on the ghost nodes.

The initial and final configurations for \mathbf{B} and for \mathbf{J} are shown in Figure 4.14. Figure 4.15 shows the computed solution for $B_\theta(r)$ at $t = 2$ relative to the analytic profile, as well as the convergence of its error. Unsurprisingly, the magnetic field and the current density profiles spread at a rate determined by the magnetic diffusivity.


The rest of the test problems in this chapter deal with moving conductors and test the implementation of the full magnetohydrodynamic algorithm. Many of these problems are ideal MHD problems that are designed with astrophysical plasmas in mind, but we have selected the problems that we find to be most pertinent to the laboratory plasmas we study in Chapter 5.

4.5.7 The advection of a magnetic pulse

One problem that plagues nascent Eulerian MHD codes is the difficulty of advecting a magnetic pulse subject to the awful $\nabla \times (\mathbf{v} \times \mathbf{B})$ term [72]. This requires an upwinding scheme that can provide enough diffusion to quell oscillations in \mathbf{B} without excessively smoothing the signal. The Flux-Constrained Transport method [33] was formulated in part as a response to this problem.


In SPH, this problem is solved trivially: we prescribe a value of \mathbf{B} to each of a set of nodes and push the nodes in the direction we want them to go. Nevertheless, it is reasonable to show that this can in fact be done. To avoid the effects of $\mathbf{J} \times \mathbf{B}$ on the surrounding medium, we use a square pulse with height $B_0 = 10^{-9}$ G in a uniform medium moving at $v = 1$ cm/s with $\rho = 100$ and $u = 1$. The results of the “trivial advection” of a pulse and its current density are shown in Figure 4.16. The sharpness of the pulse is completely unaltered by the advection. This is also evident in the profile for the current density shown in Figure 4.17.

An interesting variant of this problem is to add magnetic diffusion through artificial or physical resistivity. Price [98] did this to demonstrate his algorithm for smoothing discontinuities in \mathbf{B} . Since we are concerned with terrestrial plasmas with finite conductivities, we track the advection of a diffusing magnetic pulse using




user: jeff
Fri Feb 6 11:06:18 2009

(a)




user: jeff
Fri Feb 6 11:05:15 2009

(b)



user: jeff
Fri Feb 6 11:05:55 2009

(c)



user: jeff
Fri Feb 6 11:05:37 2009

(d)

Figure 4.14: Initial and final configurations for \mathbf{B} and \mathbf{J} in the cylindrical magnetic flux tube problem. The artifacts along the x and y axes appear both in the initial and evolved solutions and result from interpolation to a structured mesh. (a) $\mathbf{B}(x, 0)$, (b) $J_z(x, 0)$, (c) $\mathbf{B}(x, 2)$, (d) $J_z(x, 2)$.

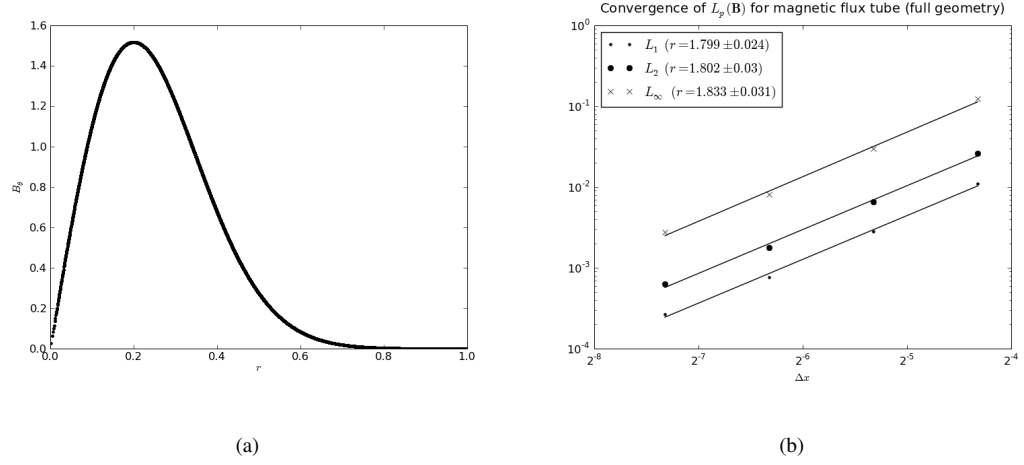


Figure 4.15: (a) Comparison of computed B_θ (dots) with exact solution (ticks) for the cylindrical magnetic flux tube problem. (b) The convergence of $L_p(\mathbf{B})$ for this problem.




Figure 4.16: Advection of a square magnetic pulse in a perfect conductor with $\nu = 1$ (a) initially, (b) halfway across the domain, and (c) after a crossing. Since SPH is Lagrangian, there is no degradation in the signal.




Figure 4.17: Current density of the advecting magnetic pulse (a) initially, (b) halfway across the domain, and (c) after a crossing.

the SPH magnetic diffusion algorithm. This test problem is nontrivial for two reasons. First, it allows us to test our implementation of periodic boundary conditions for the magnetic diffusion algorithm. Second, the accuracy achieved in this test problem is limited by the coupling between the magnetohydrodynamic and diffusion algorithms, which are integrated differently and combined with the operator-splitting technique we have discussed. The profiles for \mathbf{B} and \mathbf{J} in a conductor with diffusivity $D = 0.01$ are shown respectively in Figure 4.18 and Figure 4.19. The analytic profile is plotted with small dots, while the computed profile uses the usual large dots. Evidently, the rate of decay of the pulse is unaffected by the advection as expected. The rates of convergence of the computed \mathbf{B} and \mathbf{J} are shown in Figure 4.20.




Figure 4.18: Advection of a square magnetic pulse in a perfect conductor with $\nu = 1$ (a) initially, (b) halfway across the domain, and (c) after a crossing. The decay in the pulse is caused only by the magnetic diffusion equation.

This demonstrates that our operator splitting algorithm works as expected, and that the magnetic diffusion

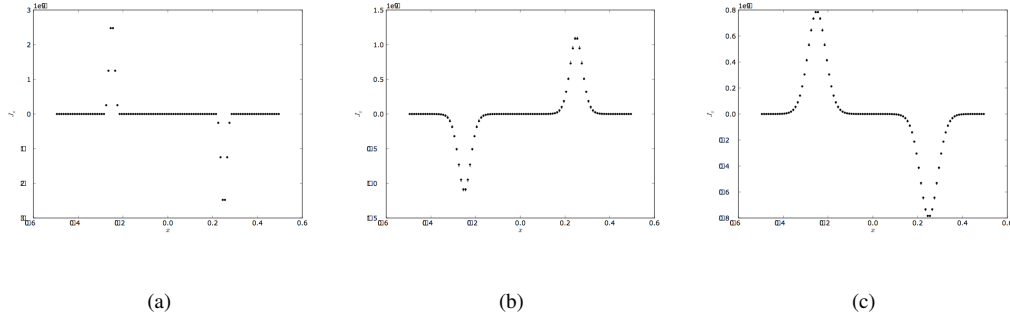


Figure 4.19: Current density of the advecting magnetic pulse (a) initially, (b) halfway across the domain, and (c) after a crossing.




Figure 4.20: Convergence of the solution for (a) \mathbf{B} and (b) \mathbf{J} in the advection of a magnetic pulse through a medium with diffusivity $D = 0.01$.

algorithm handles periodic boundary conditions correctly.

4.5.8 Magnetosonic waves

The linearized MHD equations are described in many texts ([62], [63]), and describe the evolution of magnetosonic waves and Alfvén waves. Following Price [84] again, we test magnetosonic waves in a way very similar to that in which we tested ordinary acoustic waves in Chapter 3: a wave consisting of a set of disturbances in ρ , \mathbf{v} , u , and \mathbf{B} is allowed to propagate across a periodic domain an integer number of times, and the final signal is compared to the initial. The solutions to the ideal MHD equations describing magnetosonic waves are

$$\rho(x, t) = \rho_0 + \rho_1 \sin[k(x - Vt)] \quad (4.72)$$

$$\mathbf{v}(x, t) = \mathbf{v}_1 \sin[k(x - Vt)] \quad (4.73)$$

$$p(x, t) = (\gamma - 1)\rho_0 u_0 + c_s^2 \rho_1 \sin[k(x - Vt)] \quad (4.74)$$

$$u(x, t) = \frac{p(x, t)}{(\gamma - 1)\rho(x, t)} \quad (4.75)$$

$$\mathbf{B}(x, t) = \mathbf{B}_0 + \mathbf{B}_1 \sin[k(x - Vt)] \quad (4.76)$$

where k is the wave number of the disturbance, the zero-subscripted quantities are the “ambient” values of these quantities, and the one-subscripted values are their perturbations. Here, V is the magnetosonic speed of the wave, which can be either the fast (v_+) or slow (v_-) speed as defined by (2.55). Note that \mathbf{v} and \mathbf{B} are now fully three-dimensional vectors, regardless of the symmetry of the system. The components of \mathbf{v}_1 and \mathbf{B}_1 are determined by the linearization process to be

$$v_{1x} = \frac{V\rho_1}{\rho_0} \quad (4.77)$$

$$v_{1y} = -\frac{B_{0x}B_{0y}v_{1x}}{\rho_0 W} \quad (4.78)$$

$$v_{1z} = -\frac{B_{0x}B_{0z}v_{1x}}{\rho_0 W} \quad (4.79)$$

$$B_{1x} = 0 \quad (4.80)$$

$$B_{1y} = \frac{VB_{0y}v_{1x}}{W} \quad (4.81)$$

$$B_{1z} = \frac{VB_{0z}v_{1x}}{W} \quad (4.82)$$

where we have used $W = V^2 - B_{0x}^2/\rho_0$.

For this problem, we test the propagation of a fast magnetosonic wave with $\rho_0 = 1.004$ and $\rho_1/\rho_0 = 0.0055$ and $\mathbf{B}_0 = (0.5, 0.5, 0.5)$. The mass density is initialized in the same way as it was for the acoustic wave problem. We have tempered the artificial viscosity and resistivity for this problem, using the values $\alpha = 0.1$ and $\alpha_B = 0.1$. The density disturbance is shown after 1 and 5 crossings in Figure 4.21, using the uncorrected unsymmetric energy equation with $n_h = 2$. Its convergence is shown in Figure 4.22.




Figure 4.21: Density profile for fast magnetosonic wave after (a) 1 crossing and (b) 5 crossings with $n_h = 2$ using the unsymmetric energy equation ($\alpha = \alpha_B = 0.1$)

The convergence rate of the solution is noticeably lower than that of the pure hydrodynamic case—

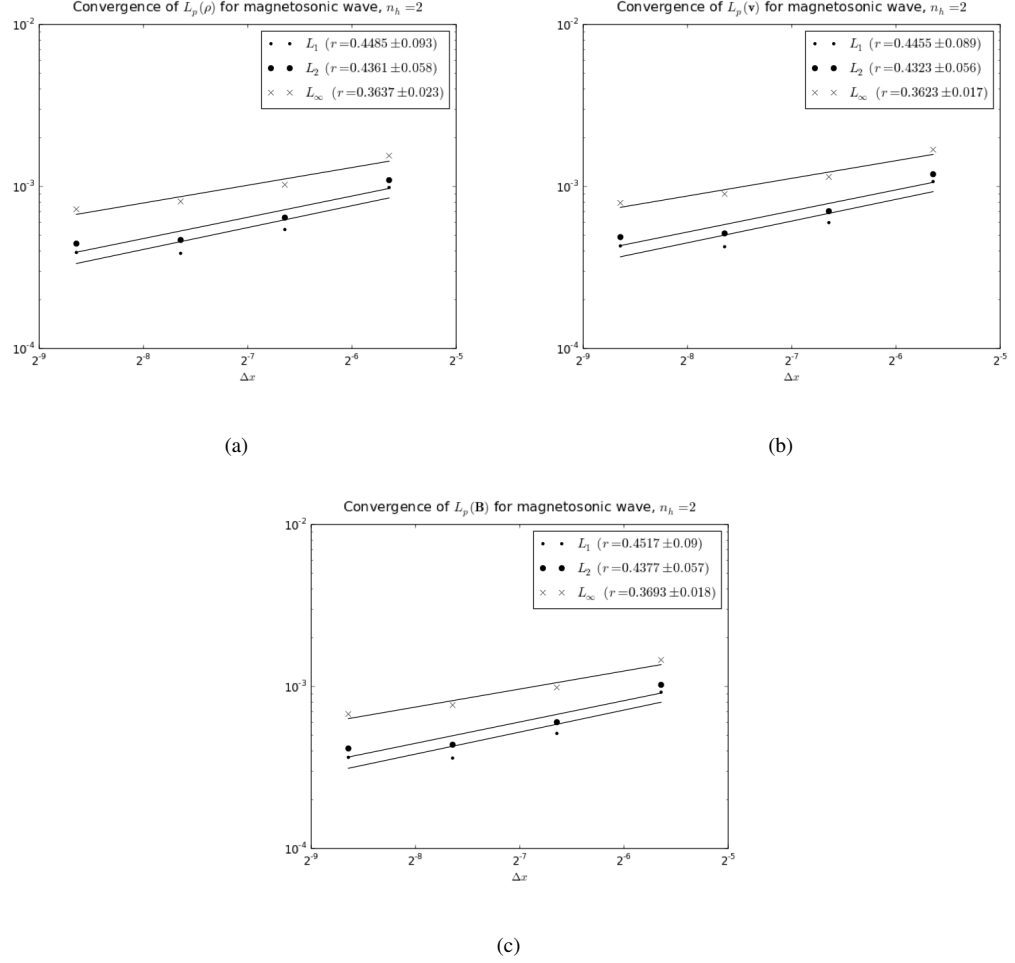


Figure 4.22: Convergence of solutions for (a) ρ (b) v , and (c) B for fast magnetosonic wave after 5 crossings with $n_h = 2$ using the unsymmetric energy equation ($\alpha = \alpha_B = 0.1$)

less than 0.5! This is a trend we will notice in SPMHD as we explore more test problems. If no artificial dissipation is used, the simulation will complete (unlike in the hydrodynamic case), but the convergence is negative—the solution gets worse with higher resolution. Table 4.1, Table 4.2, and Table 4.3 show these rates of convergence for various values of n_h . When compared to the corresponding acoustic wave tables Table 3.4 and Table 3.5, it is plain that the convergence of SPMHD for magnetosonic waves is far worse than that of SPH for acoustic waves. The solution for the mass density does not appear to converge for any value of n_h less than 2! $n_h = 1.6$ remains a curiously bad value for all quantities, and $n_h = 1.8$ fares little better.

n_h	$r\{L_1(\rho)\}$	$r\{L_2(\rho)\}$	$r\{L_\infty(\rho)\}$
1.2	0.01282 ± 0.00016	-0.000776 ± 0.00012	$-0.03826 \pm 4.3e-5$
1.4	$-0.0002275 \pm 4.3e-9$	$0.003708 \pm 1.8e-5$	-0.008448 ± 0.0034
1.6	-0.05784 ± 0.00049	-0.03449 ± 0.00024	-0.01606 ± 0.00015
1.8	0.03855 ± 0.0044	0.02192 ± 0.0026	$0.02282 \pm 1.9e-5$
2.0	0.4485 ± 0.093	0.4361 ± 0.058	0.3637 ± 0.023

Table 4.1: Convergence rates for L_p norms of the mass density in the magnetosonic wave as functions of n_h .

n_h	$r\{L_1(\mathbf{v})\}$	$r\{L_2(\mathbf{v})\}$	$r\{L_\infty(\mathbf{v})\}$
1.2	0.01456 ± 0.00053	0.004983 ± 0.00046	-0.05144 ± 0.00013
1.4	0.221 ± 0.021	0.1618 ± 0.025	0.13 ± 0.0049
1.6	-0.05613 ± 0.00078	-0.06141 ± 0.00078	-0.1115 ± 0.0013
1.8	0.05514 ± 0.0021	0.0224 ± 0.0031	-0.01142 ± 0.0013
2.0	0.4455 ± 0.089	0.4323 ± 0.056	0.3623 ± 0.017

Table 4.2: Convergence rates for L_p norms of the velocity in the fast magnetosonic wave as functions of n_h .

n_h	$r\{L_1(\mathbf{B})\}$	$r\{L_2(\mathbf{B})\}$	$r\{L_\infty(\mathbf{B})\}$
1.2	0.01206 ± 0.00043	0.002211 ± 0.00037	$-0.5663 \pm 8.1e-5$
1.4	0.2209 ± 0.021	0.161 ± 0.024	0.1281 ± 0.0054
1.6	-0.06106 ± -0.00099	-0.06607 ± 0.00096	-0.1169 ± 0.0015
1.8	0.05413 ± 0.0022	0.02239 ± 0.0034	-0.0138 ± 0.001
2.0	0.4517 ± 0.09	0.4377 ± 0.057	0.3693 ± 0.018

Table 4.3: Convergence rates for L_p norms of the magnetic field in the fast magnetosonic wave as functions of n_h .

It is worth studying the effects of the compatible energy method and the ∇h corrections on the

solution to this problem in order to see whether the convergence rate can be improved. The convergence for ρ , \mathbf{v} and \mathbf{B} using the ∇h terms is shown in Figure 4.23. Evidently, these corrections have no discernible effect on the convergence of this problem.

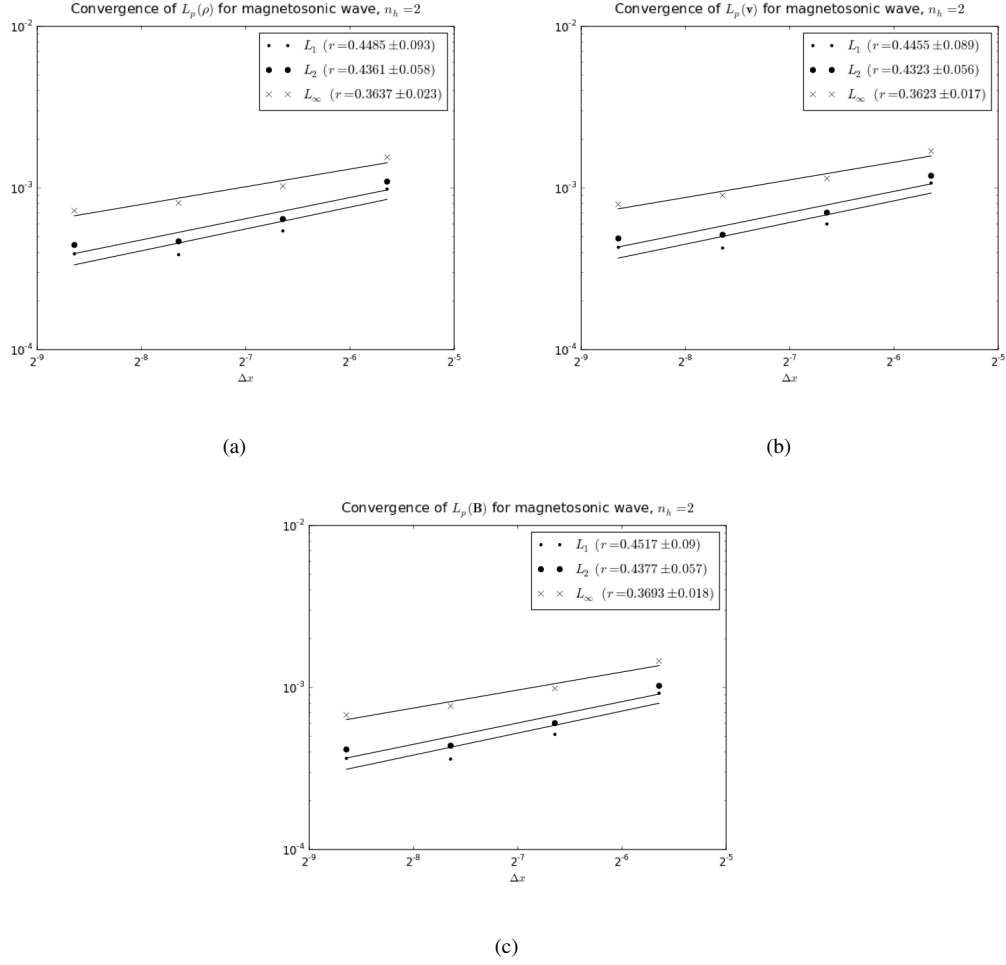


Figure 4.23: Convergence of solutions for (a) ρ , (b) \mathbf{v} , and (c) \mathbf{B} for fast magnetosonic wave after 5 crossings with $n_h = 2$ using the ∇h corrections to the unsymmetric energy equation ($\alpha = \alpha_B = 0.1$)

The convergence for these quantities using the compatible energy method is shown in Figure 4.24. The convergence of the L_p error norms for ρ is significantly improved, but at the price of a minor degradation in those for the other quantities.

In any case, it is clear that while our algorithm produces convergent solutions for this problem, it

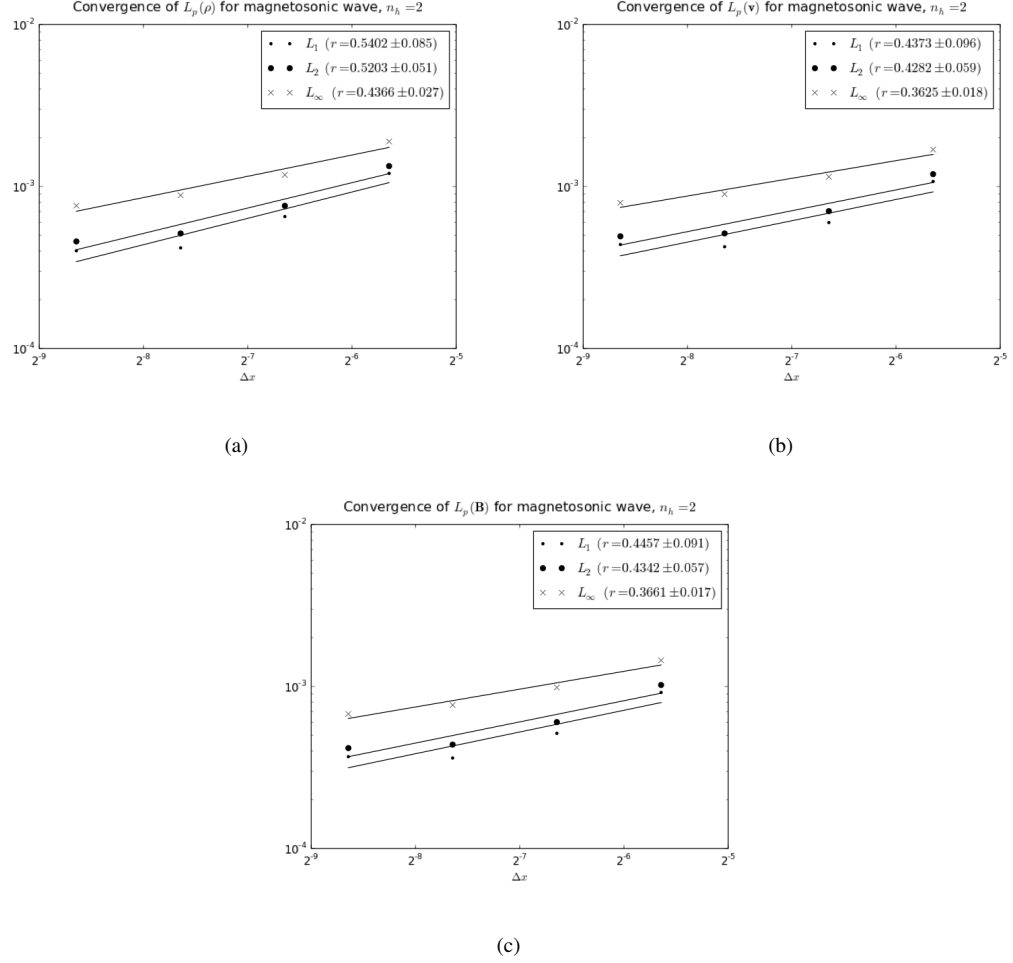


Figure 4.24: Convergence of solutions for (a) ρ (b) v and (c) B for fast magnetosonic wave after 5 crossings with $n_h = 2$ using the compatible energy equation ($\alpha = \alpha_B = 0.1$)

is not on a par with SPH’s treatment of pure acoustic waves. Nevertheless, we find that the n_h parameter continues to play an important role in the convergence of the solution.

4.5.9 Alfvén waves

It is equally important to ensure that our algorithm can handle Alfvén waves, which are those purely magnetic disturbances caused by the “tension” ($\mathbf{B} \cdot \nabla \mathbf{B}$) term in the magnetic stress tensor. We now test the propagation of a circularly-polarized Alfvén wave across a 1D periodic domain in a manner described by Tóth [26].

The wave travels through a medium of uniform density ($\rho_0 = u_0 = 1$) and is initially of the form

$$\mathbf{v}(x, t) = (0, 0.1 \sin(2\pi x), 0.1 \cos(2\pi x)) \quad (4.83)$$

$$\mathbf{B}(x, t) = (1, 0.1 \sin(2\pi x), 0.1 \cos(2\pi x)) \quad (4.84)$$

in numeric units in which $\mu_0 = 1$. To within 1%, the Alfvén speed $v_A = 1$. Thus, we expect the wave to propagate without distortion across a domain of unit length in the same way as the weak magnetosonic waves we have previously described. Since there are no discontinuities in \mathbf{B} , we set the artificial resistivity parameter α_B to 0. Figure 4.25 shows B_y for the Alfvén wave after 1 and 5 crossings. The wave is compared to its initial configuration after 5 crossings, and the convergence of the error in \mathbf{v} and \mathbf{B} using the unsymmetric energy equation with $n_h = 2$ is shown in Figure 4.26. Pleasingly, the convergence rate for both approaches the expected value of 2 for continuous flows. This suggests that the low convergence rates we saw for magnetosonic waves stem from the method’s treatment of the magnetic pressure term in the Lorentz Force, and not necessarily from the evolution of the magnetic field itself.

However, this convergence is once again shown to be dependent on n_h . Table 4.4 and Table 4.5 show that decreasing n_h severely damages the convergence, destroying it completely below $n_h = 1.4$. This non-convergence results from large phase errors in the wave when not enough neighbors are sampled. In fact, only $n_h = 2$ comes anywhere near the desired convergence rate.




Figure 4.25: Solution for B_y (a) after 1 crossing and (b) after 5 crossings for circularly-polarized Alfvén wave with $n_h = 2$ using the unsymmetric energy equation




Figure 4.26: Convergence of solutions for (a) \mathbf{v} and (b) \mathbf{B} for circularly-polarized Alfvén wave after 5 crossings with $n_h = 2$ using the unsymmetric energy equation

n_h	$r\{L_1(\mathbf{v})\}$	$r\{L_2(\mathbf{v})\}$	$r\{L_\infty(\mathbf{v})\}$
1.2	-0.1129 ± 0.0077	-0.1129 ± 0.0077	-0.1129 ± 0.0077
1.4	0.07641 ± 0.17	0.07641 ± 0.17	0.0761 ± 0.17
1.6	0.05428 ± 0.0023	0.05428 ± 0.0023	0.05428 ± 0.0023
1.8	0.2527 ± 0.037	0.2527 ± 0.037	0.2527 ± 0.037
2.0	1.959 ± 0.0013	1.959 ± 0.0013	1.959 ± 0.0013

Table 4.4: Convergence rates for L_p norms of \mathbf{v} in the 1D Alfvén wave as functions of n_h .

n_h	$r\{L_1(\mathbf{B})\}$	$r\{L_2(\mathbf{B})\}$	$r\{L_\infty(\mathbf{B})\}$
1.2	-0.111 ± 0.0081	-0.111 ± 0.0081	-0.111 ± 0.0081
1.4	0.1483 ± 0.26	0.1483 ± 0.26	0.1483 ± 0.26
1.6	0.05889 ± 0.0018	0.05889 ± 0.0018	0.05889 ± 0.0018
1.8	0.2761 ± 0.033	0.2761 ± 0.033	0.2761 ± 0.033
2.0	1.867 ± 0.0083	1.857 ± 0.0083	1.867 ± 0.0083

Table 4.5: Convergence rates for L_p norms of \mathbf{B} in the 1D Alfvén wave as functions of n_h .

Examples of the phase errors that erode the convergence of the errors in the 1D Alfvén wave problem for different values of n_h are shown in Figure 4.27.

4.5.10 Brio and Wu’s shock tube

Magnetized shock tube problems have become an important tool for demonstrating the effectiveness of MHD algorithms ([32], [36], [67], [68], [69]). In these problems, two conducting gases with respective initial states $(\rho_1, \mathbf{v}_1, u_1, \mathbf{B}_1)$ and $(\rho_2, \mathbf{v}_2, u_2, \mathbf{B}_2)$ are separated by a partition, and this partition is lifted, after which a set of magnetosonic shocks and rarefactions develop. If \mathbf{B} is tangentially discontinuous at the contact discontinuity, that discontinuity is propagated as an Alfvén wave.

The shock conditions of these problems are complicated, and analytic solutions have not typically been written down. However, recently-developed Godunov methods based on the numerical solution for the MHD shock conditions have given physically plausible answers to these problems, and these answers are used to apprise other MHD algorithms. In the last chapter, we described a method for comparing our solutions to the Sod problem to solutions computed by a Godunov code. We use this method to evaluate the answers we compute for MHD shock tubes, and other problems where possible.

In particular, a test problem designed by Brio and Wu [32] has become the standard shock tube problem for MHD in the same way that the Sod problem has for ordinary hydrodynamics. Stationary ideal gas with $\gamma = 2$ fills a shock tube. The initial states for the left and right halves of the tube are

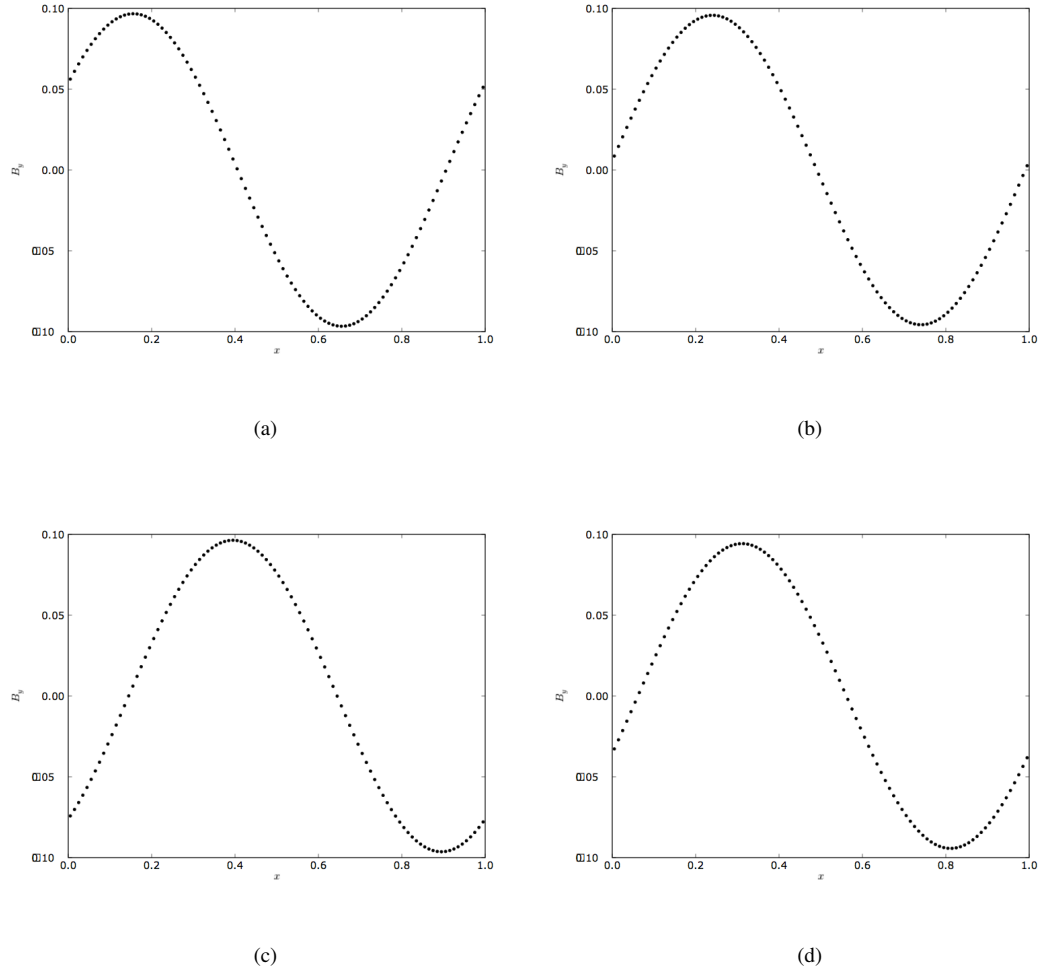


Figure 4.27: Alfvén wave after 5 crossings with (a) $n_h = 1.2$, (b) $n_h = 1.4$, (c) $n_h = 1.6$, and (d) $n_h = 1.8$ using the uncorrected unsymmetric energy equation

$$\begin{aligned}
\rho_1 &= 1 & \rho_2 &= 0.125 \\
p_1 &= 1 & p_2 &= 0.1
\end{aligned} \tag{4.85}$$

$$\begin{aligned}
\mathbf{B}_1 &= (0.75, 1, 0) & \mathbf{B}_2 &= (0.75, -1, 0)
\end{aligned} \tag{4.86}$$

where, again, we have used numeric units of measure in which $\mu_0 = 1$. We proceed in the same way we solved the Sod problem, studying the convergence of the problem using the unsymmetric energy equation with n_h assuming various values between 1.2 and 2.6. We use rigid boundaries at the ends of the shock tube, which we place at $x = \pm 0.5$. To treat discontinuities in \mathbf{v} and \mathbf{B} , artificial viscosity and resistivity are used with $\alpha = \alpha_B = 1$, following Price. We address the issue of the tensile instability by removing the components of magnetic stress related to B_x , since these components cannot have dynamic effects in 1D. We then compare the computed solutions at $t = 0.1$ to those obtained by Athena on a 3200-cell grid—these solutions are the solid lines behind the dotted SPH solutions in the figures. The rates of convergence for the L_1 error norms of the solution using the unsymmetric energy equation are tabulated in Table 4.6. This table gives us an idea of how the average error decreases when more resolution is added and as n_h is increased.

The convergence rates are much lower in general than those we attained for the Sod problem, and increasing n_h only improves things very slowly. In fact, the convergence rate for \mathbf{v} flattens out around $n_h = 1.6$ and even begins to suffer at high values of this parameter. This is because the tensile instability begins to have a significant effect as the stencil grows beyond this threshold. To illustrate this effect, we have plotted curves for ρ in Figure 4.28 and v_x in Figure 4.29 that show the appearance and growth of this instability as n_h is increased. The convergence is limited for low n_h by the insufficient sampling of neighbors, and for high n_h by the emergence of the tensile instability.

These plots show a number of undesirable features that do not improve as n_h is increased. The solutions at $n_h = 1.6$ are shown in Figure 4.30 for reference. We attempt to address these features with the same refinements to the energy equation that we have used before. The compatible energy method with

n_h	$r\{L_1(\rho)\}$	$r\{L_1(\mathbf{v})\}$	$r\{L_1(u)\}$	$r\{L_1(p)\}$	$r\{L_1(\mathbf{B})\}$
1.2	0.4971 ± 0.009	0.4388 ± 0.00075	0.3951 ± 0.0073	0.4093 ± 0.0073	0.4583 ± 0.00068
1.3	0.4943 ± 0.0059	0.4396 ± 0.00093	0.4095 ± 0.0079	0.4127 ± 0.0057	0.4472 ± 0.00054
1.4	0.5066 ± 0.0029	0.4605 ± 0.00044	0.4408 ± 0.008	0.4465 ± 0.0039	0.4571 ± 0.00057
1.5	0.5079 ± 0.0091	0.4781 ± 0.00023	0.4719 ± 0.0081	0.4638 ± 0.0055	0.4636 ± 0.00025
1.6	0.5194 ± 0.0015	0.4946 ± 0.00016	0.4966 ± 0.0076	0.4826 ± 0.006	0.4767 ± 0.00028
1.7	0.5408 ± 0.0015	0.5054 ± 0.00011	0.5123 ± 0.0072	0.5008 ± 0.0058	0.5083 ± 0.00017
1.8	0.5511 ± 0.0017	0.5066 ± 0.00011	0.5185 ± 0.0069	0.5111 ± 0.004	$0.5323 \pm 5e-5$
1.9	0.5541 ± 0.0021	$0.5026 \pm 5.5e-5$	0.5201 ± 0.006	0.5123 ± 0.0027	$0.5421 \pm 1.9e-5$
2.0	0.555 ± 0.0019	$0.4968 \pm 2e-5$	0.5207 ± 0.005	0.5111 ± 0.0021	$0.5411 \pm 5.3e-5$
2.1	0.5593 ± 0.0014	$0.4939 \pm 2.6e-5$	0.5246 ± 0.0041	0.5151 ± 0.0019	0.5428 ± 0.00018
2.2	0.5635 ± 0.001	$0.4936 \pm 7.7e-5$	0.5285 ± 0.0034	0.5181 ± 0.0018	0.5479 ± 0.00042
2.3	0.569 ± 0.00074	0.4945 ± 0.00015	0.5323 ± 0.0028	0.5224 ± 0.0017	0.5544 ± 0.00085
2.4	0.5742 ± 0.00052	0.497 ± 0.0002	0.5372 ± 0.0022	0.5314 ± 0.0019	0.5622 ± 0.0015
2.5	0.5772 ± 0.00043	0.497 ± 0.00024	0.5423 ± 0.0016	0.5448 ± 0.0024	0.5688 ± 0.0024
2.6	0.5808 ± 0.00053	0.4953 ± 0.00025	0.5453 ± 0.0012	0.5587 ± 0.0031	0.5726 ± 0.0036

Table 4.6: Convergence rates for L_1 error norms of the solution for the Brio/Wu shock tube as a function of n_h . The unsymmetric energy equation was used to produce these results.




Figure 4.28: Onset of the tensile instability at $t = 0.1$ in the mass density: (a) $n_h = 1.6$, (b) $n_h = 1.8$. Note the clumping of the nodes around the contact discontinuity.




Figure 4.29: Onset of the tensile instability at $t = 0.1$ in the velocity: (a) $n_h = 1.6$, (c) $n_h = 1.8$. Slight oscillations start to appear to the right of the contact discontinuity.

artificial thermal conductivity using Morris's variable switch produces the L_1 convergence rates shown in Table 4.7; the curves for $n_h = 1.6$ are shown in Figure 4.31. The ∇h terms are added to produce the rates in Table 4.8, and their curves for $n_h = 1.6$ are shown in Figure 4.32. These methods generally improve the accuracy of the solution and work much more effectively than simply increasing n_h . For example, the compatible energy method brings the velocity, specific thermal energy, pressure, and \mathbf{B} curves much closer to Athena's around the contact discontinuity, resulting in significantly higher convergence rates. Likewise, the thermal conduction term stabilizes the jump in u in the same region, though the variable switch reduces the conduction too severely and allows some of the overshoot to remain. The ∇h correction helps some quantities and hurts others, but seems to increase the sensitivity of the instability to higher values of n_h . In any case, these methods do nothing to stop the instability. The velocity convergence as shown in Table 4.7 and Table 4.8 is the best indication that the instability is still present. Thus, we are not able to apply the simple rules of thumb that seemed to work for our hydrodynamic simulations. A larger n_h is not always better. We must choose a good stencil size that does not excite the tensile instability— $n_h = 1.6$, say—if we are to use the algorithm without significant modification.

The improved convergence rates are significantly higher than those of the uncorrected algorithm,

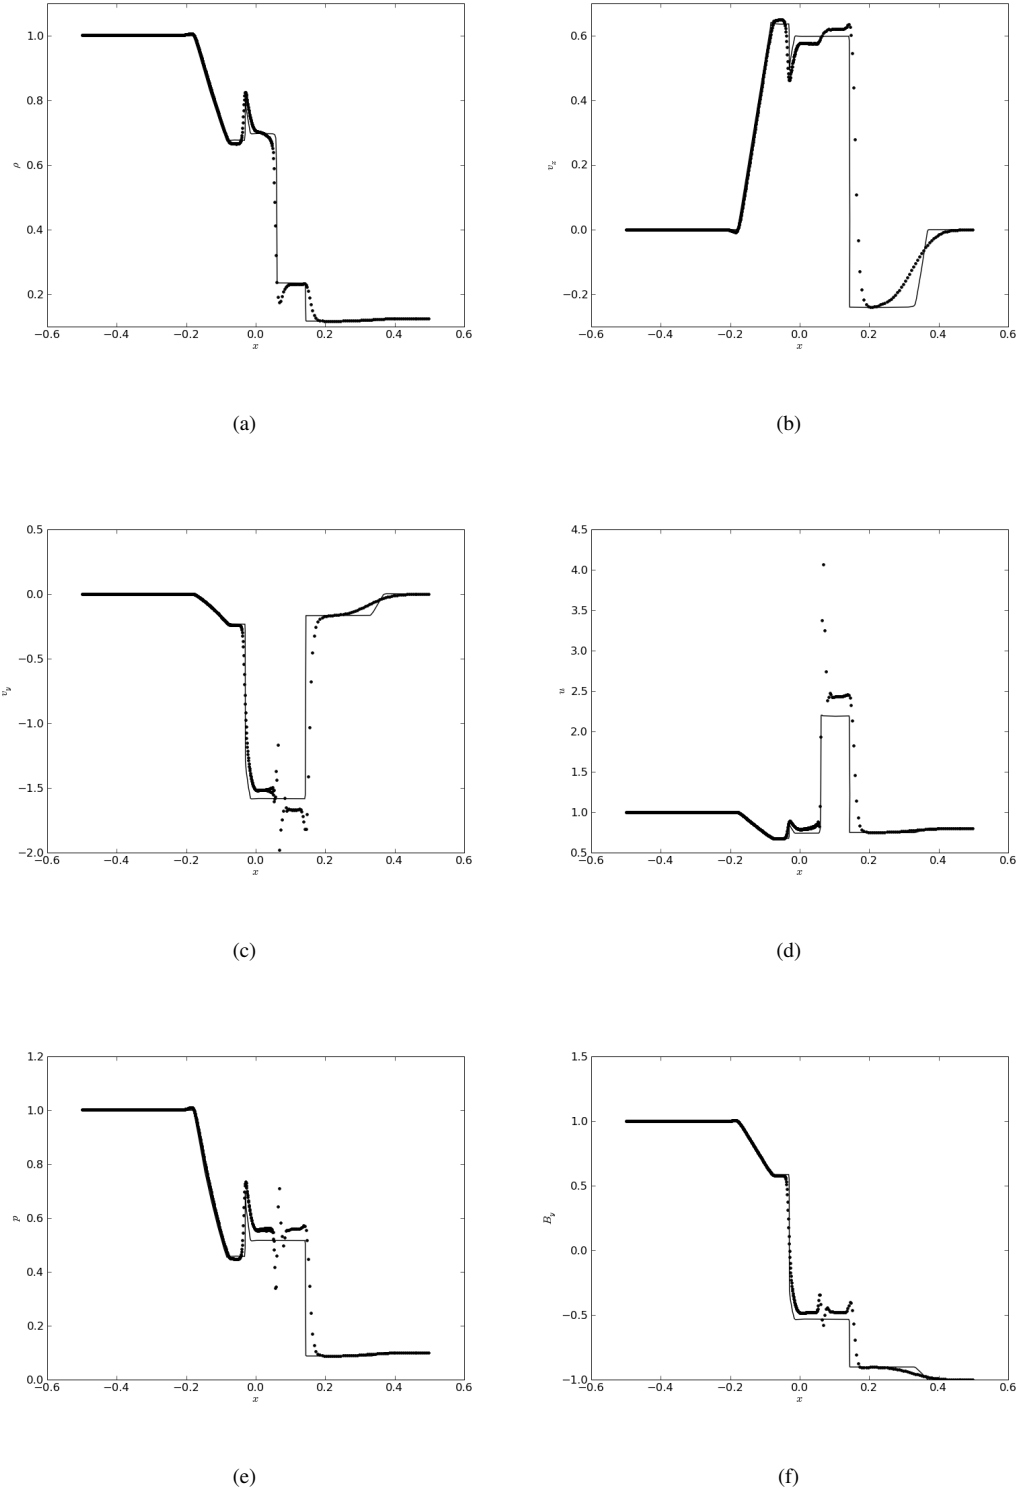


Figure 4.30: Solutions to the Brio/Wu shock tube problem ($n_h = 1.6$) at $t = 0.1$ using the unsymmetric energy equation: (a) ρ , (b) v_x , (c) v_y , (d) u , (e) p , (f) B_y .

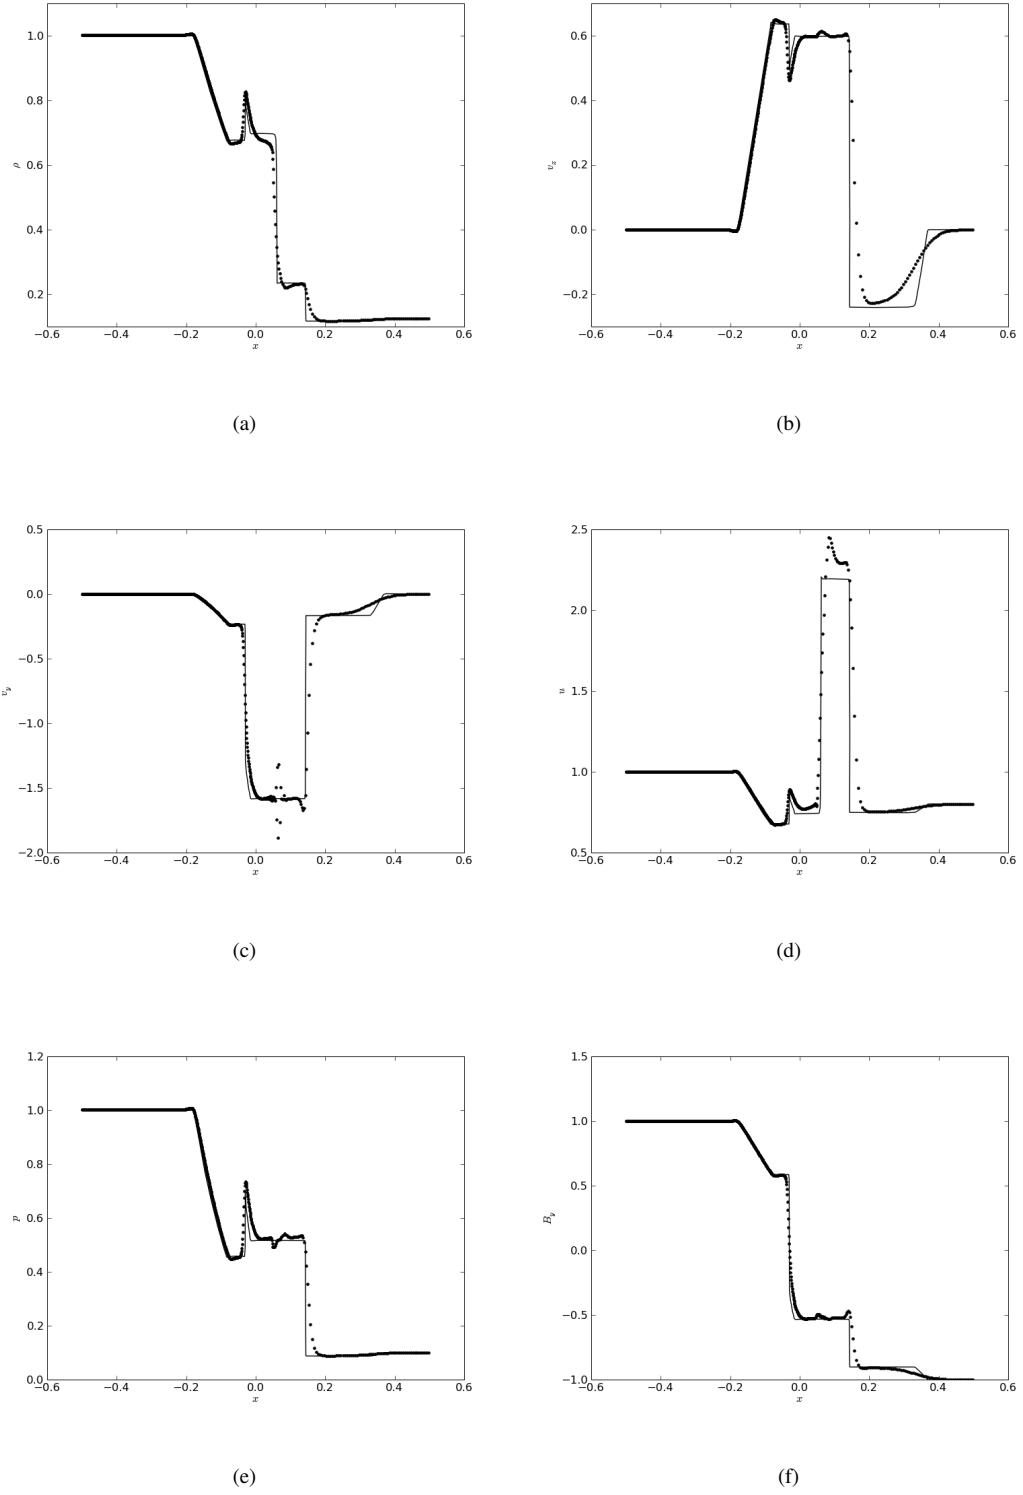


Figure 4.31: Solutions to the Brio/Wu shock tube problem ($n_h = 1.6$) at $t = 0.1$ using the compatible energy method and artificial thermal conduction with a variable α_u : (a) ρ , (b) v_x , (c) v_y , (d) u , (e) p , (f) B_y .

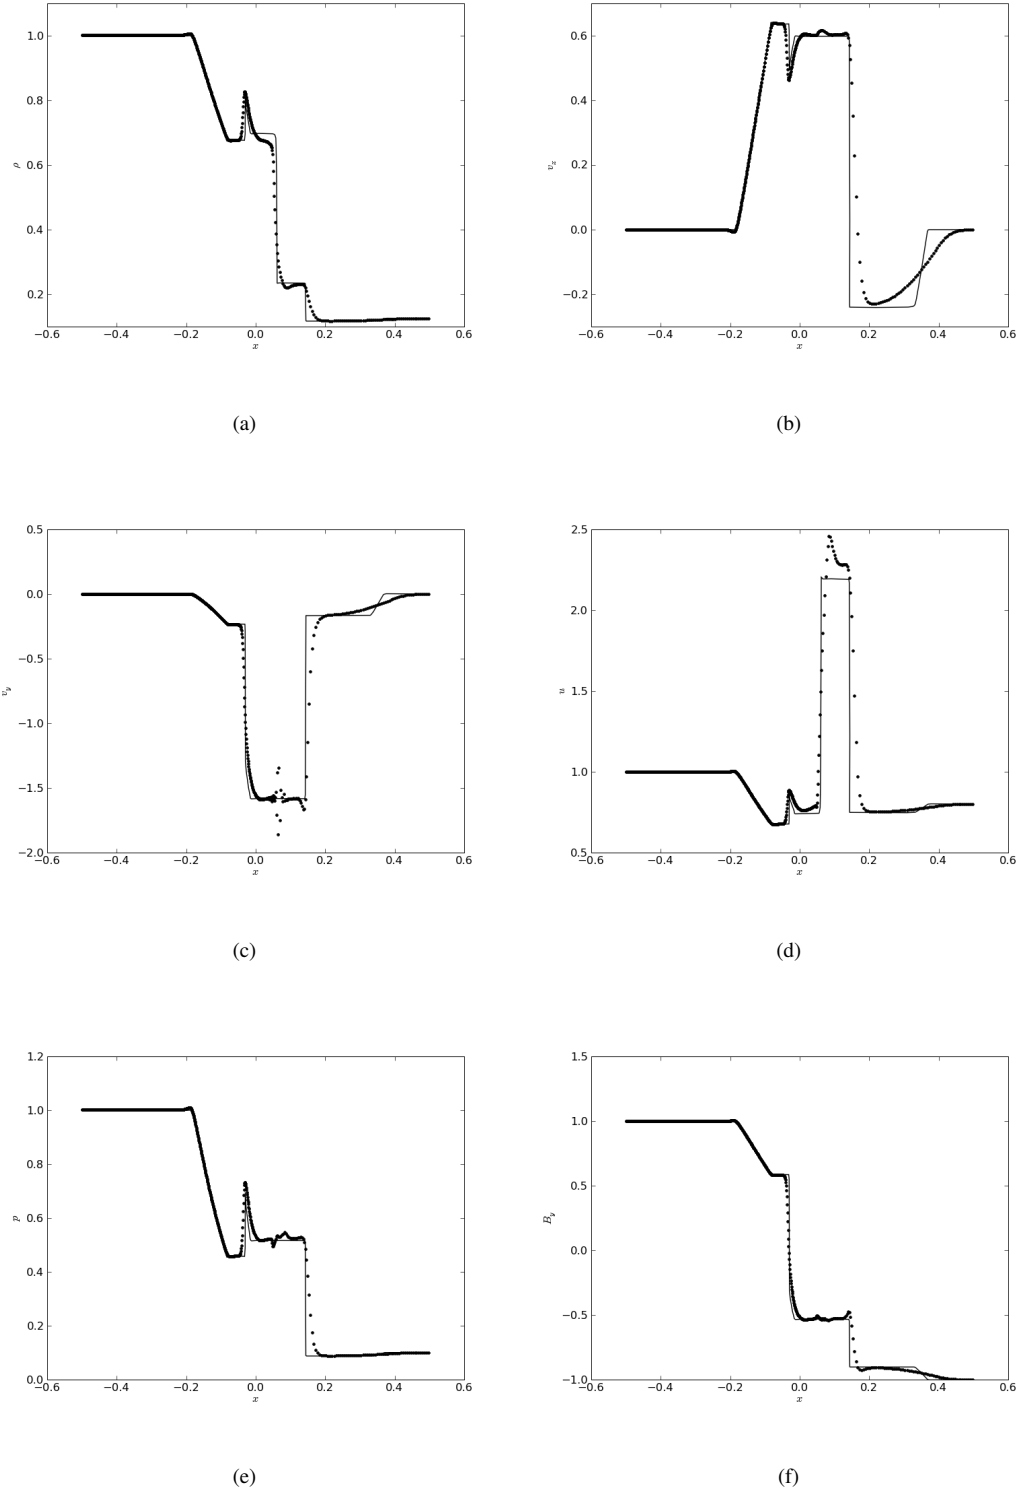


Figure 4.32: Solutions to the Brio/Wu shock tube problem ($n_h = 1.6$) at $t = 0.1$ using the compatible energy method, artificial thermal conduction (variable α_u), and ∇h corrections: (a) ρ , (b) v_x , (c) v_y , (d) u , (e) p , (f) B_y .

n_h	$r\{L_1(\rho)\}$	$r\{L_1(\mathbf{v})\}$	$r\{L_1(u)\}$	$r\{L_1(p)\}$	$r\{L_1(\mathbf{B})\}$
1.2	0.5544 ± 0.00083	0.5917 ± 0.0016	0.6867 ± 0.001	0.643 ± 0.0013	0.6367 ± 0.00025
1.4	0.5313 ± 0.001	0.6076 ± 0.0037	0.7049 ± 0.00048	0.6605 ± 0.00081	0.6358 ± 0.0006
1.6	0.5312 ± 0.00077	0.6071 ± 0.0058	0.718 ± 0.0012	0.6423 ± 0.00056	$0.6381 \pm 1.1e - 5$
1.8	0.5592 ± 0.0012	0.624 ± 0.0075	0.7222 ± 0.0017	0.6751 ± 0.0013	$0.7126 \pm 3.9e - 5$
2.0	0.5453 ± 0.0025	0.5818 ± 0.0048	0.6794 ± 0.002	0.649 ± 0.0018	0.7012 ± 0.0006

Table 4.7: Convergence rates for L_1 error norms of the solution for the Brio/Wu shock tube as a function of n_h . These results were produced using the compatible energy method and artificial thermal conduction with Morris's variable switch α_u .

n_h	$r\{L_1(\rho)\}$	$r\{L_1(\mathbf{v})\}$	$r\{L_1(u)\}$	$r\{L_1(p)\}$	$r\{L_1(\mathbf{B})\}$
1.2	0.5723 ± 0.0052	0.6363 ± 0.0061	0.7132 ± 0.00059	0.711 ± 0.00038	0.654 ± 0.0002
1.4	0.5299 ± 0.0028	0.6169 ± 0.0059	0.7158 ± 0.00044	0.6885 ± 0.00052	$0.6206 \pm 7.3e - 5$
1.6	0.5405 ± 0.0034	0.6154 ± 0.0092	0.7219 ± 0.00053	0.6988 ± 0.0013	0.6561 ± 0.00013
1.8	0.5533 ± 0.0034	0.5998 ± 0.01	0.7105 ± 0.0011	0.699 ± 0.0025	0.6789 ± 0.00044
2.0	0.543 ± 0.0027	0.5871 ± 0.0094	0.7026 ± 0.0025	0.686 ± 0.0045	0.69 ± 0.00049

Table 4.8: Convergence rates for L_1 error norms of the solution for the Brio/Wu shock tube as a function of n_h . These results were produced using the compatible energy method, artificial thermal conduction (with a variable α_u), and ∇h corrections.

but SPMHD still has a long way to go before it achieves its potential. The accuracy of the magnetic field and the velocity can be improved further by tempering the artificial dissipation, but one must be careful to ensure that oscillations are kept to a minimum. For example, our use of Morris's variable switch α_u for artificial thermal conduction is too severe in its limiting, and the residual overshoot in u produces oscillations in other quantities around the contact discontinuity. On the other hand, we saw in the Sod problem in Chapter 3 that setting $\alpha_u = 1$ does get rid of these oscillations but significantly erodes the convergence of u . Better limiters for artificial dissipation in \mathbf{v} , \mathbf{B} , and u would help this algorithm tremendously.

4.5.11 Planar magnetized compression

A shock tube problem explored by Dai and Woodward [36] involves the collision of two magnetized gases. This interests us because we explore the compression of magnetized plasma in the next chapter. The initial states for the two interacting gases in the shock tube are

$$\begin{array}{ll} \rho_1 = 1 & \rho_2 = 1 \\ p_1 = 20 & p_2 = 1 \end{array} \quad (4.87)$$

$$\mathbf{v}_1 = (1, 0, 0) \quad \mathbf{v}_2 = (-1, 0, 0) \quad (4.88)$$

$$\mathbf{B}_1 = (1.4105, 1.4105, 0) \quad \mathbf{B}_2 = (1.4105, 1.4105, 0) \quad (4.89)$$

with numeric units in which $\mu_0 = 1$. We run this problem both in 1D and 2D to demonstrate our multidimensional algorithm. We use inflow boundary conditions at $x = -0.5$ and $x = 0.5$ to prevent dynamics near the boundaries. The initial node distributions in our study have 50, 100, and 200 nodes, but these numbers quickly grow as nodes flow into the problem. Unlike Brio and Wu's problem, this simulation tends to yield better results at $n_h = 2$ than at $n_h = 1.6$. In fact, we observe more clumping at $n_h = 1.6$ than at $n_h = 2$ in this case! The conditions that determine the best value of n_h are quite problem-dependent. For this problem, we choose $n_h = 2$. The solution at $t = 0.08$ using the unsymmetric energy equation and a variable switch for artificial thermal conductivity is shown in Figure 4.33, and the convergence is shown in Figure 4.34. The specific thermal energy profile develops large oscillations at the origin if artificial thermal conductivity is not used.

The solution for the magnetic field is not convergent at all— B_z clearly plateaus at the wrong values on the left and right sides. v_z has similar problems, missing its left and right plateaus in directions opposite to the deviations of B_y . The solution appears to conserve energy to some degree, but the balance of kinetic and magnetic energy is incorrect. If we use the compatible energy method, the convergence rates are somewhat improved for \mathbf{v} and \mathbf{B} , but the problem essentially remains. The ∇h correction factor seems to address this issue effectively. Figure 4.35 shows the velocity profiles for the compatible energy method and ∇h corrections and illustrates the effects of the latter on v_z , while Figure 4.36 shows the effects of these two improvements on v_z .

We have mentioned previously that the “compatible” energy method as adapted to MHD does not perfectly

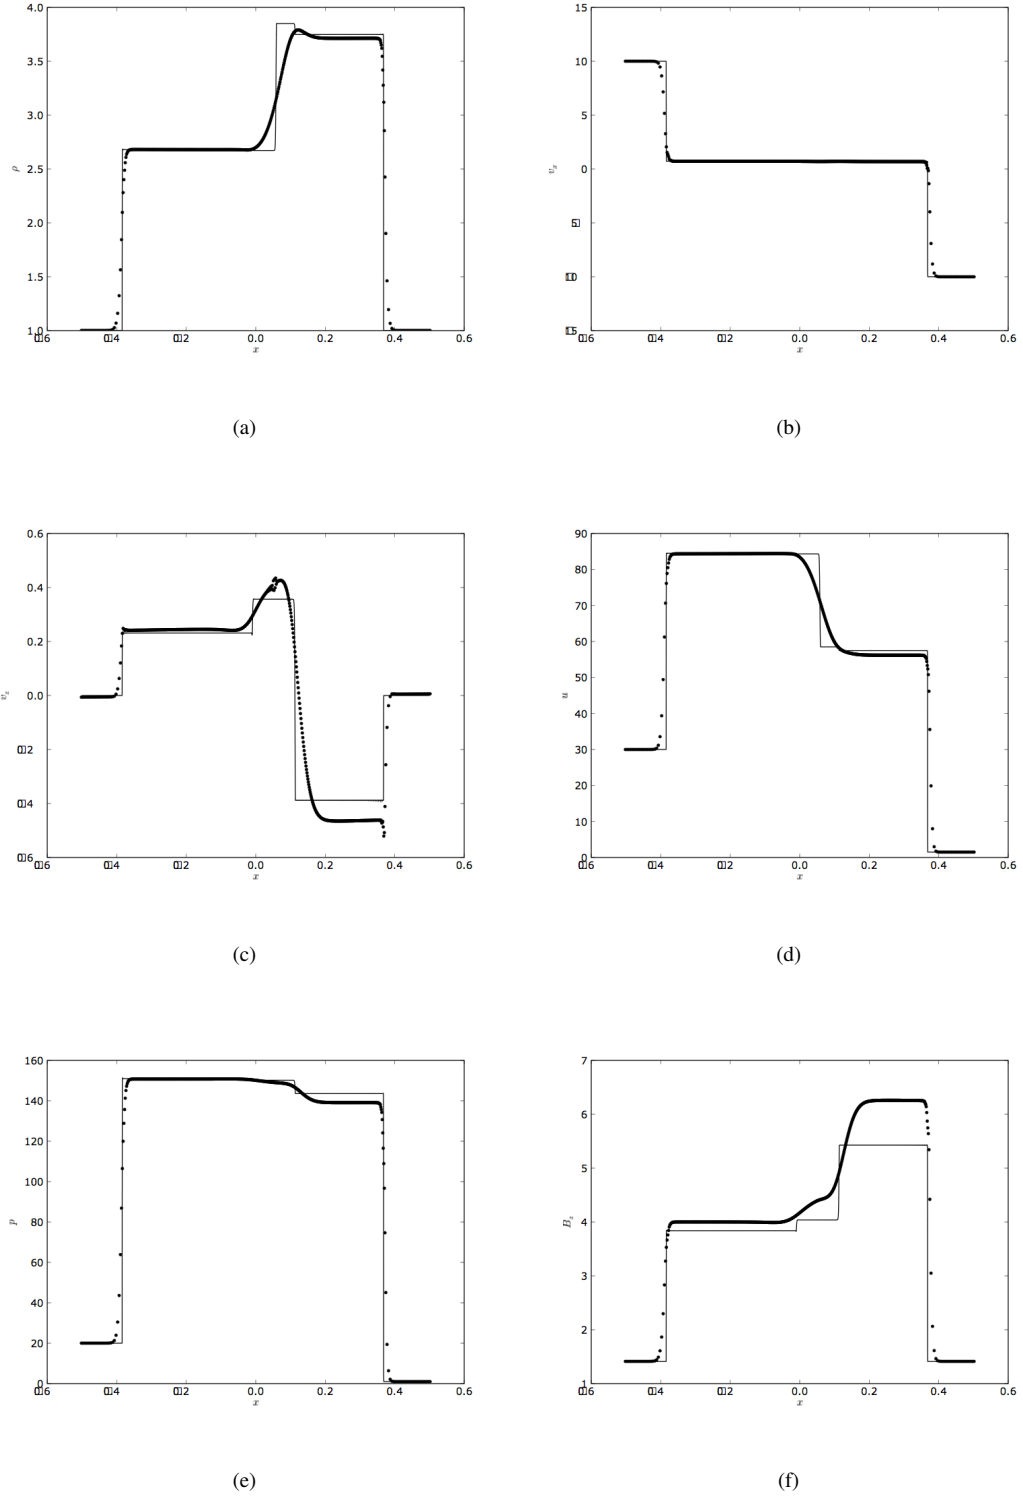


Figure 4.33: Solutions to the magnetic compression problem ($n_h = 2$) at $t = 0.08$ using the unsymmetric energy equation: (a) ρ , (b) v_x , (c) v_z , (d) u , (e) p , (f) B_z .

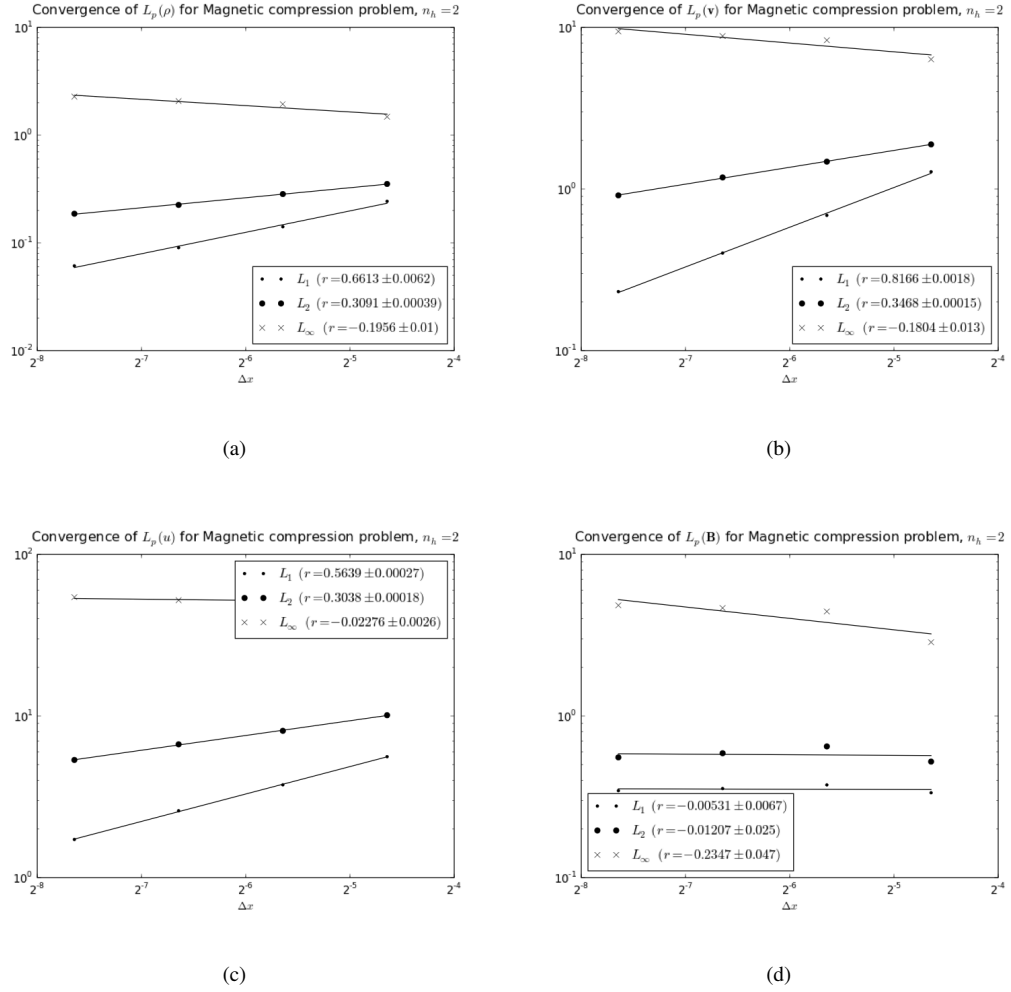


Figure 4.34: Convergence rates for the L_p norms of the solutions to the magnetic compression problem ($n_h = 2$) using the unsymmetric energy equation: (a) ρ , (b) \mathbf{v} , (c) u , (d) \mathbf{B} .

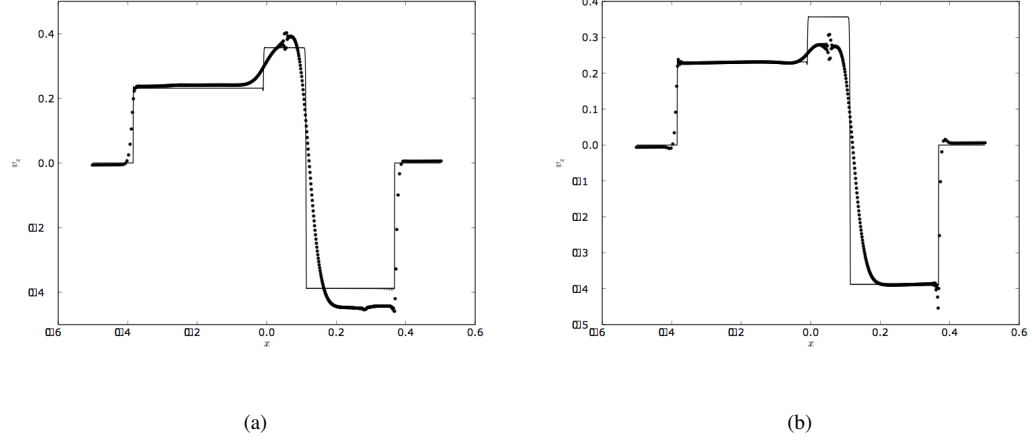


Figure 4.35: Solution for v_z in the planar magnetic compression problem using (a) the compatible energy method and (b) the ∇h corrections.




Figure 4.36: Solution for B_z in the planar magnetic compression problem using (a) the compatible energy method and (b) the ∇h corrections.

conserve energy because of the incompatible differencing used to evaluate $\frac{D\mathbf{v}}{Dt}$ and $\frac{D\mathbf{B}}{Dt}$. If we were able to evolve these two variables consistently, it is likely that these results would improve dramatically. In the meantime, incorporating the ∇h correction into these disparate evolution equations seems to gain us enough accuracy to achieve a passable solution. Interestingly, the compatible energy method does get a better solution for v_z near the contact discontinuity. The convergence rates for the L_1 norms of the solutions using the unsymmetric energy equation, the compatible energy method, the ∇h correction, and the combination of the two improvements are listed for comparison in Table 4.9. This combination produces significantly improved results as compared to the original unsymmetric energy equation. In particular, it appears that the two improvements are complementary for this problem, each producing improvements in different regions of the problem.

Method	$r\{L_1(\rho)\}$	$r\{L_1(\mathbf{v})\}$	$r\{L_1(u)\}$	$r\{L_1(\mathbf{B})\}$
Uncorrected	0.6613 ± 0.0062	0.8166 ± 0.0018	0.5639 ± 0.00027	-0.00531 ± 0.0067
Compatible	0.7708 ± 0.023	0.9241 ± 0.017	0.9112 ± 0.053	0.1543 ± 0.0016
∇h	0.7257 ± 0.0014	0.896 ± 0.0006	0.709 ± 0.0038	0.7115 ± 0.0042
Compatible + ∇h	0.8012 ± 0.013	0.9659 ± 0.01	0.9775 ± 0.037	0.7616 ± 0.0016

Table 4.9: Convergence rates for L_1 error norms of the solution for the magnetic compression problem using the different methods for the energy equations.

4.5.12 Magnetic rarefaction

The next problem treats the rarefaction waves caused by an expanding magnetized gas in a shock tube, with the left and right sides moving away from each other. The initial conditions are

$$\begin{aligned}
\rho_1 &= 1 & \rho_2 &= 1 \\
p_1 &= 1 & p_2 &= 1 \\
\mathbf{v}_1 &= (-1, 0, 0) & \mathbf{v}_2 &= (1, 0, 0) \\
\mathbf{B}_1 &= (1, 0, 0) & \mathbf{B}_2 &= (1, 0, 0),
\end{aligned} \tag{4.90}$$

$$\tag{4.91}$$

again using our numeric units with $\mu_0 = 1$. We use outflow boundary conditions at the ends of the shock tube $x = -0.5$ and $x = 0.5$ to prevent unwanted behavior. Initially, the node distributions contain 100, 200, 400, and 800 nodes. Like the last problem, we observe better results in 1D at $n_h = 2$ than at $n_h = 1.6$, so we continue to use this value. The solution at $t = 0.1$ using the unsymmetric energy equation with a variable artificial thermal conductivity is shown in Figure 4.37, and the convergence is illustrated in Figure 4.38. The transverse components of \mathbf{v} and \mathbf{B} remain essentially zero, so they are not depicted.

The convergence rates for the uncorrected unsymmetric energy equation are actually admirable for this problem. The most undesirable feature is the jump in each quantity at the rarefaction front. None of our improvements are able to alleviate this. The other oddity is a kink in the density and energy profiles that begins 2/3 of the way to the contact discontinuity from the rarefaction front. The ∇h corrections are able to mitigate this kink in both quantities. By contrast, the compatible energy method does not address this kink, and produces a larger dip at the contact discontinuity in an attempt to resolve it, reducing the accuracy of the solution. Even combining the compatible energy method with the ∇h terms does not alleviate the effects of this increased error. The profiles for the mass density obtained using the two improvements are shown in Figure 4.39. The convergence rates for the L_1 norms of ρ , \mathbf{v} , u , and p using the various methods are shown in Table 4.10.

It is not surprising that the ∇h terms make a large difference in the solution to this problem. The smoothing scale h is proportional to the density, which varies linearly over a large portion of the problem. The kink in the density profile is probably caused by the change in the values of h from the neighbor contributions to the nodes in that region, since it occurs at a location at which the contact discontinuity becomes invisible

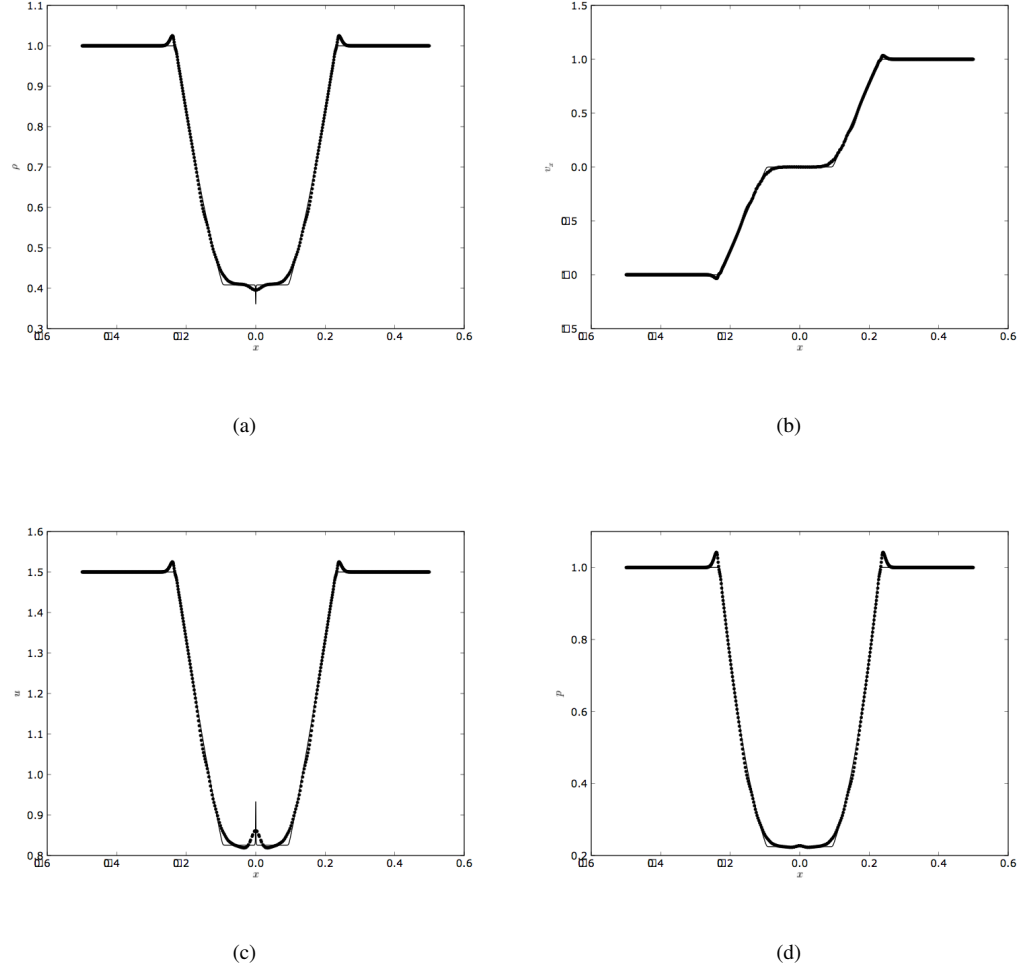


Figure 4.37: Solutions to the magnetic rarefaction problem ($n_h = 2$) at $t = 0.1$ using the unsymmetric energy equation: (a) ρ , (b) v_x , (c) u , (d) p .

Method	$r\{L_1(\rho)\}$	$r\{L_1(\mathbf{v})\}$	$r\{L_1(u)\}$	$r\{L_1(p)\}$
Uncorrected	0.7725 ± 0.011	0.8364 ± 0.017	0.919 ± 0.007	0.891 ± 0.0047
Compatible	0.6755 ± 0.026	0.8262 ± 0.021	0.969 ± 0.0066	0.9287 ± 0.00049
∇h	0.806 ± 0.0038	0.8474 ± 0.0082	0.9242 ± 0.0053	0.8965 ± 0.0047
Compatible + ∇h	0.7412 ± 0.0074	0.8413 ± 0.011	0.9649 ± 0.0048	0.9211 ± 0.0026

Table 4.10: Convergence rates for L_1 error norms of the solution for the magnetic rarefaction problem using the different methods for the energy equations.

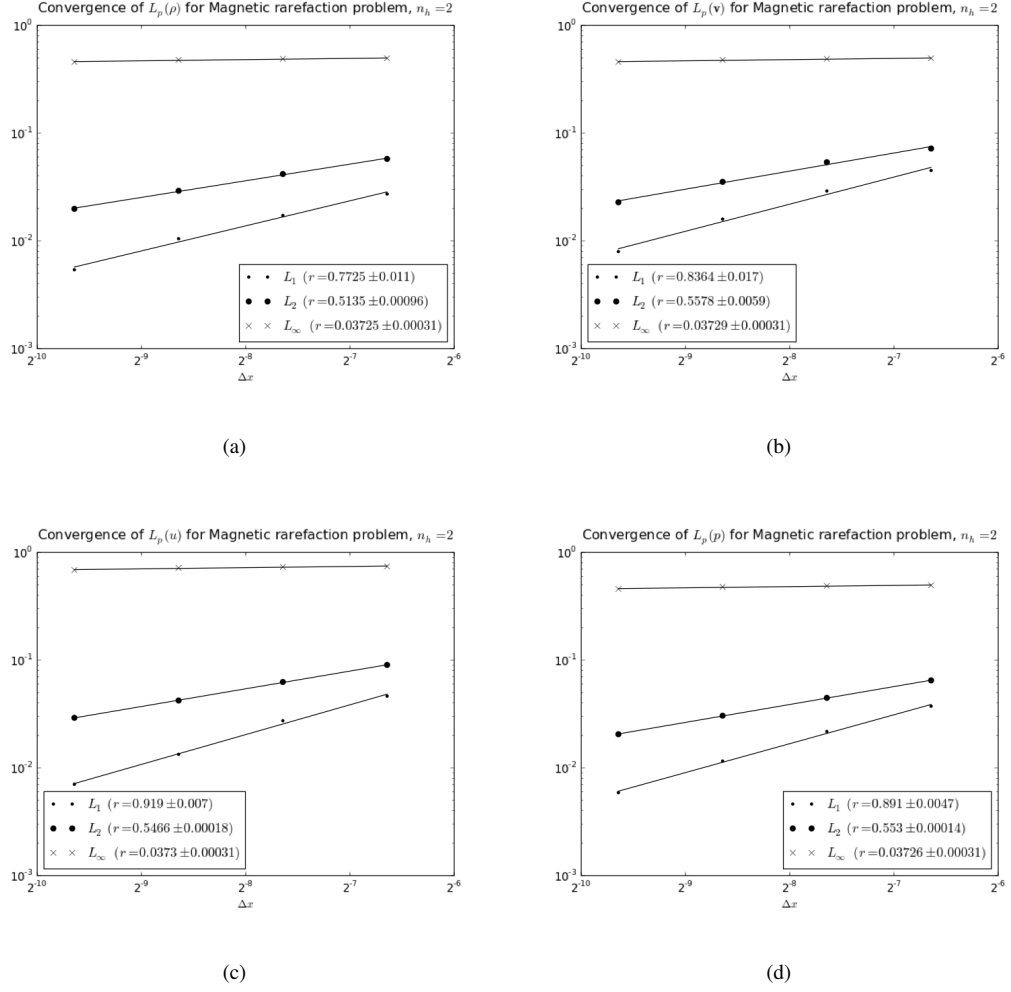


Figure 4.38: Convergence rates for the L_p norms of the solutions to the magnetic compression problem ($n_h = 2$) using the unsymmetric energy equation: (a) ρ , (b) \mathbf{v} , (c) u , (d) p .




Figure 4.39: Solution for ρ in the magnetic rarefaction problem using (a) the compatible energy method and (b) the ∇h corrections.

to the rest of the gas. The compatible energy method tries to resolve the discontinuity (which appears both in ρ and u) but the evacuation of nodes from this region severely limits its resolution, and the kink is likely made worse, resulting in a loss of accuracy. This is a good example of how a qualitative improvement (the attempted resolution of a discontinuity) can result in degraded accuracy.

4.5.13 Spherical Blast Wave

The multidimensional propagation of magnetosonic waves in a magnetized medium can be studied using a problem described by Balsara [139]. A stationary uniform adiabatic ($\gamma = 1.4$) gas with density $\rho = 1$ is seeded with energy such that the pressure $p = 1000$ within a spherical region of radius $r = 0.05$ about the origin and $p = 1$ elsewhere. The gas is magnetized with a uniform 10 G magnetic field. We model this problem in 2D in the $\theta = \pi/2$ plane after Price and several others using a uniform $N \times N$ lattice of nodes, where N assumes the values 50, 100, 200, and 400. We use the stencil parameter $n_h = 2$. As in our earlier MHD simulations, we use a variable switch α_u enabling artificial thermal conduction, and the usual settings for artificial viscosity. We use a variable switch α_B to control the artificial diffusion of the magnetic field. The boundaries in this problem are insignificant, as no signal reaches them in the allotted simulation time. We use

periodic boundary conditions in x and y . Our solution at $t = 0.02$ is compared with that computed by Athena on a 400×400 grid of uniform cells. Figure 4.40 shows the solutions rendered by Athena and our method for ρ and $|\mathbf{B}|$, side by side.

The SPH answer is interpolated using the SPH approximation to Athena's grid, and L_p error norms are computed from the difference in the two solutions. In this case, the best results are achieved by combining the ∇h terms and the compatible energy method; Figure 4.41 shows the convergence rates achieved using this method.


Table 4.11 shows the convergence rates of the L_1 norms for the errors in the various quantities computed using our myriad methods. These rates are consistent with the preceeding MHD results in that the ∇h terms improve the answer significantly. Ironically, the uncorrected energy equation yields results that converge to those of Athena at the highest rate.

Method	$r\{L_1(\rho)\}$	$r\{L_1(\mathbf{v})\}$	$r\{L_1(u)\}$	$r\{L_1(p)\}$
Uncorrected	0.3977 ± 0.029	0.2825 ± 0.0082	0.4921 ± 0.0018	0.4501 ± 0.16
Compatible	0.6304 ± 0.029	0.4011 ± 0.0014	0.4418 ± 0.0031	0.6871 ± 0.024
∇h	0.5524 ± 0.0039	0.3069 ± 0.0056	0.4542 ± 0.0016	0.4911 ± 0.025
Compatible + ∇h	0.7181 ± 0.018	0.4114 ± 0.0022	0.4234 ± 0.0019	0.6606 ± 0.0094


Table 4.11: Convergence rates for L_1 error norms of the solution for Balsara's blast wave using the different methods for the energy equations.

The question arises as to whether Athena's solutions converge at the expected rates. Figure 4.42 shows the convergence of Athena's solution to itself over the resolutions $N = 50$, $N = 100$, $N = 200$, and $N = 400$, where N is the number of cells on a side. The lower resolutions are computed and interpolated to the $N = 400$ grid where they are compared with the answer itself. Though the linear interpolation scheme we use is crude, it does confirm that Athena self-converges at around first order.


Clearly, the SPH method and Athena's Godunov method converge to a similar answer. As with the other MHD tests, the rates of convergence we have measured measure the relative convergence of these two methods. Increasing the resolution of Athena's solution may give a more accurate estimate of a proper convergence rate for the SPH simulations, provided that we can trust this high-resolution solution.




(a)



(b)



(c)



(d)

Figure 4.40: Solutions to the Balsara blast wave at $t = 0.02$: (a) ρ (Athena), (b) ρ (SPH), (c) $|\mathbf{B}|$ (Athena), (d) $|\mathbf{B}|$ (SPH).

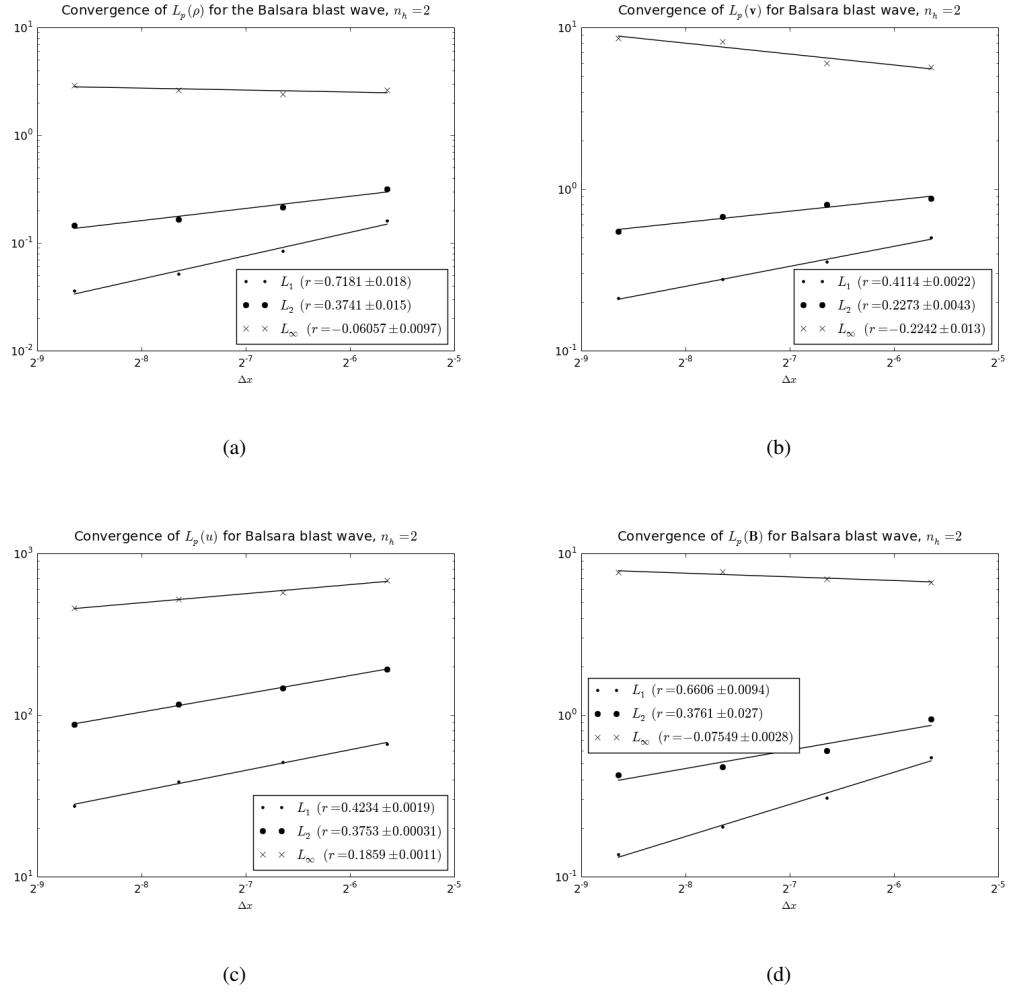


Figure 4.41: Convergence rates for the L_p norms of the Balsara blast wave at $t = 0.02$ using the ∇h corrections and the compatible energy method: (a) ρ , (b) v , (c) u , (d) B .

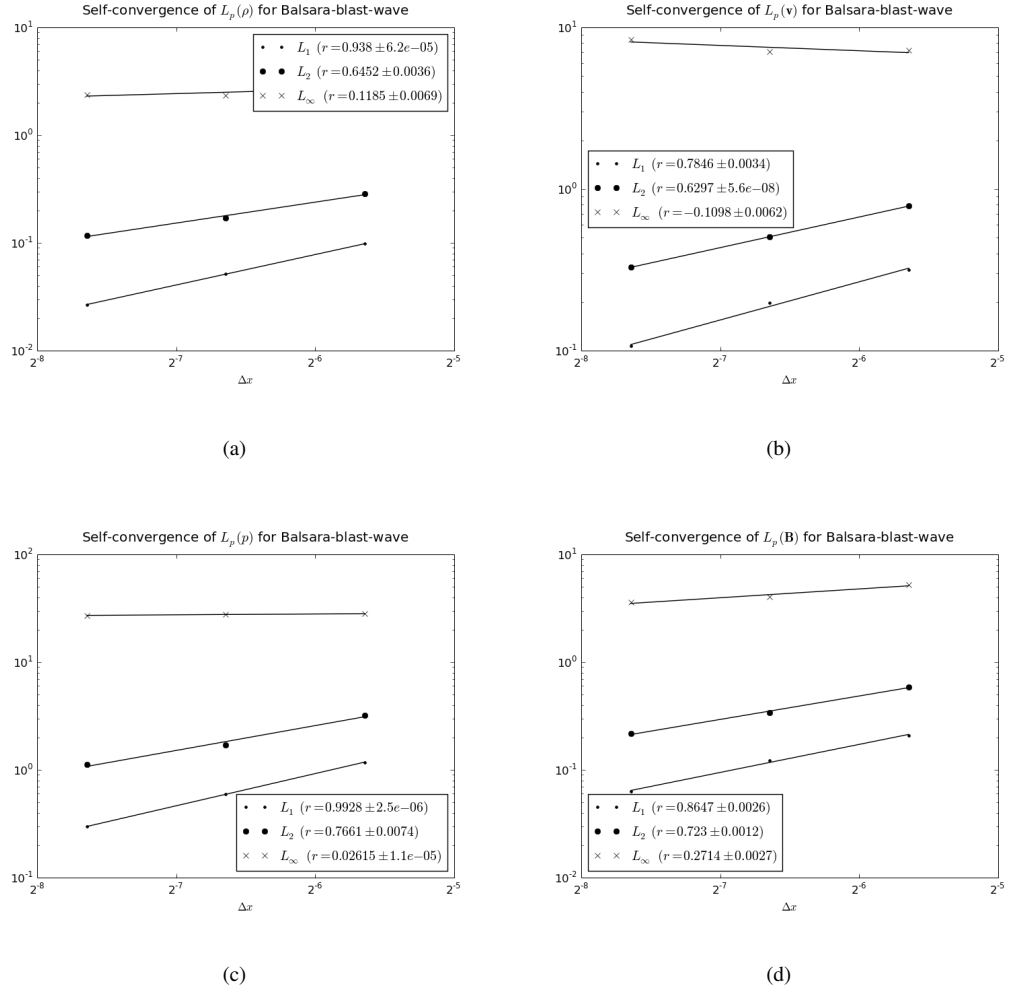


Figure 4.42: Self-convergence rates for the L_p norms of Athena's solutions to the Balsara blast wave at $t = 0.02$: (a) ρ , (b) \mathbf{v} , (c) p , (d) \mathbf{B} .

4.5.14 MHD Turbulence: the Orzsag-Tang vortex

Our last multidimensional problem, another staple of MHD test suites, is a magnetized vortex used by Orzsag and Tang [140] to study the development of magnetic microturbulence. We run a variant of this problem with a magnetized ideal gas with $\gamma = 5/3$ on a domain $[0, 1] \otimes [0, 1]$ with the initial conditions

$$\rho(x, y, 0) = \frac{25}{36\pi} \quad (4.92)$$

$$\mathbf{v}(x, y, 0) = (-\sin(2\pi y), \sin(2\pi x), 0) \quad (4.93)$$

$$u(x, y, 0) = 1 \quad (4.94)$$


$$\mathbf{B}(x, y, 0) = \frac{1}{4\pi}(-\sin(2\pi y), \sin(4\pi x), 0) \quad (4.95)$$

in a numeric system of units for which $\mu_0 = 1$. The boundaries are periodic in x and y . Once again, we compare our computed results to a reference solution computed using Athena on a uniform grid on with 400×400 cells. The simulation time at which we perform the comparison is $t = 0.5$. Our convergence study uses initially uniform lattices of nodes numbering 50, 100, 200, and 400 nodes per side.

Run with parameters similar to those for our other multidimensional tests (variable α_u , α_B , $n_h = 2$), the SPH results we obtain converge to those for Athena, albeit very slowly. The density and magnetic field configurations are shown for Athena and for the highest resolution SPH run in Figure 4.43. The L_1 error norms for the solutions appear in Table 4.12. None of the energy refinements seem to significantly improve the results.


For reference, Athena's self-convergence is given in Figure 4.44.

One explanation for the slow convergence is that the shearing present in the problem produces excessive dissipation in the velocity and magnetic field. Since we have focused on compressional problems more than rotational problems in this thesis, we have not implemented mechanisms to remove this excess dissipation. The profiles in Figure 4.43 seem to reinforce this idea—the structure is the same, but the SPH data are lower, indicating more dissipation than appears in the Riemann solver solution. Tuning the artificial dissipation to ignore shear flows would greatly improve the convergence of our method in this problem.




User: jeff
 Mon Feb 16 17:48:33 2009

(a)




User: jeff
 Mon Feb 16 17:50:02 2009

(b)



User: jeff
 Mon Feb 16 17:49:09 2009

(c)



User: jeff
 Mon Feb 16 17:50:25 2009

(d)

Figure 4.43: Solutions to the Orszag-Tang vortex at $t = 0.5$: (a) ρ (Athena), (b) ρ (SPH), (c) $|\mathbf{B}|$ (Athena), (d) $|\mathbf{B}|$ (SPH).

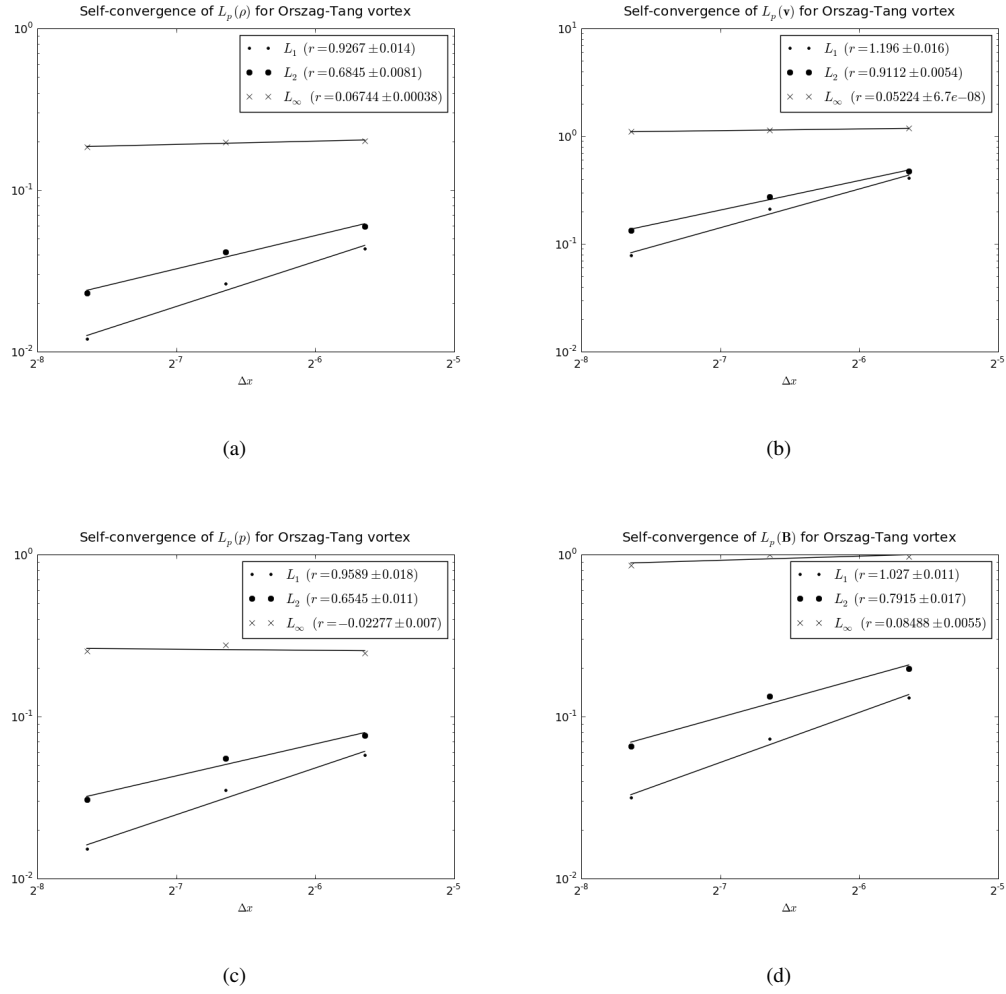


Figure 4.44: Self-convergence rates for the L_p norms of Athena's solutions to the Orszag-Tang vortex at $t = 0.5$: (a) ρ , (b) \mathbf{v} , (c) p , (d) \mathbf{B} .

Method	$r\{L_1(\rho)\}$	$r\{L_1(\mathbf{v})\}$	$r\{L_1(u)\}$	$r\{L_1(p)\}$
Uncorrected	0.3049 ± 0.00069	0.2804 ± 0.0004	0.3728 ± 0.0014	0.2962 ± 0.00022
Compatible	0.3098 ± 0.00022	0.2791 ± 0.0036	0.3784 ± 0.00076	0.2902 ± 0.00027
∇h	0.3148 ± 0.0026	0.2734 ± 0.00032	0.3962 ± 0.0013	0.2854 ± 0.00021
Compatible + ∇h	0.3208 ± 0.0034	$0.2738 \pm 2.3e-5$	0.406 ± 0.00044	$0.2801 \pm 8.2e-5$

Table 4.12: Convergence rates for L_1 error norms of the solution for the Orszag-Tang vortex using the different methods for the energy equations.

We have plotted the convergence of the L_p error norms for the uncorrected energy equation in Figure 4.45.

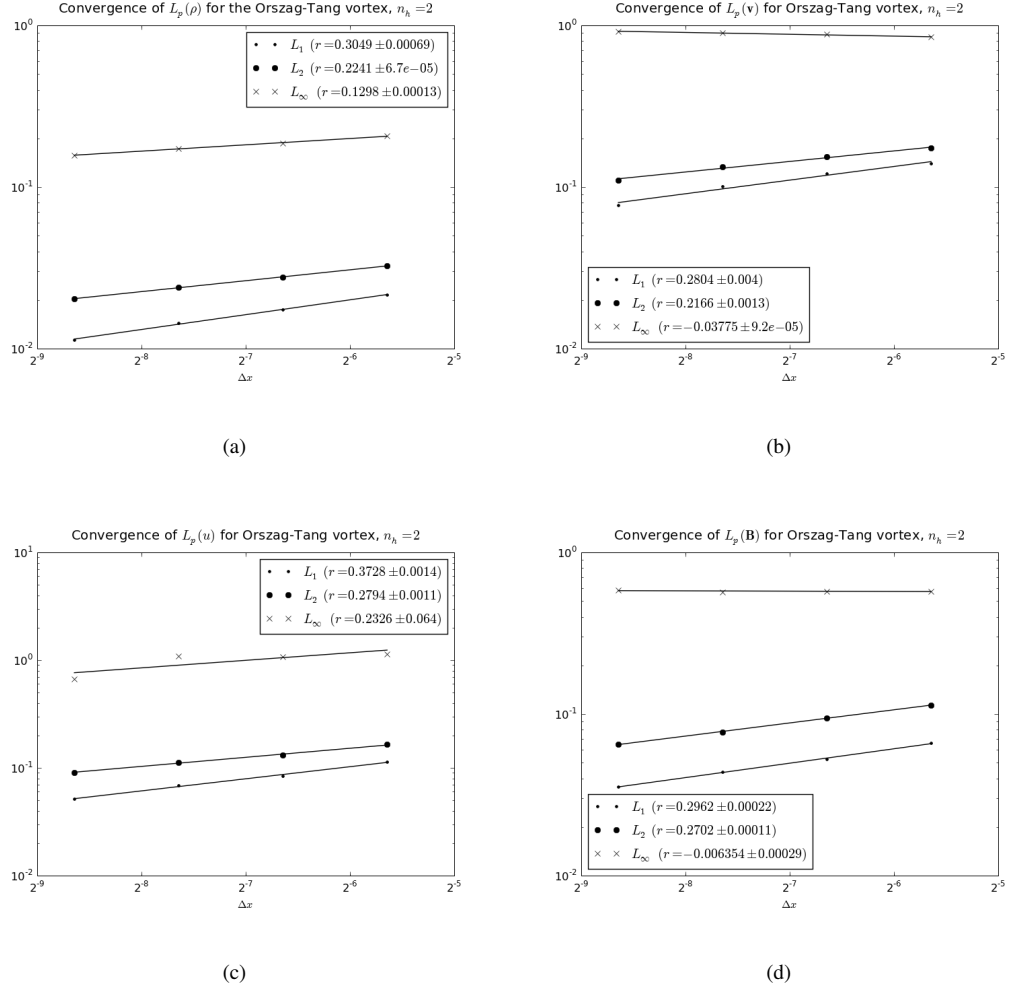


Figure 4.45: Convergence rates for the L_p norms of the Orszag-Tang vortex at $t = 0.5$ using the uncorrected energy equation: (a) ρ , (b) \mathbf{v} , (c) u , (d) \mathbf{B} .

This problem represents a good benchmark for multidimensional MHD schemes with vorticity that clearly needs further study. The use of the convergence study here underscores the difference between compressional and vortical problems.

4.6 Summary

At this point we have validated our method against various problems in magnetic diffusion and MHD. The performance of SPH for some magnetized problems is degraded by the onset of the tensile instability for values of n_h that yielded optimal results in pure hydrodynamic simulations. However, we have seen that the solutions continue to converge slowly even as this parameter is increased. This suggests to us that the instability results from problems with the discretization of $\mathbf{J} \times \mathbf{B}$ and not explicitly from these higher values of n_h . The success of the algorithm in hydrodynamics also supports this theory. However, the variation in the optimal value of n_h for the problems in this chapter makes it more difficult to make the general sort of arguments we made in the last chapter.

We have also seen that the corrections from ∇h play a much more significant role in the presence of magnetic fields than they did in pure hydrodynamics problems, even in 1D simulations. This suggests that the discretization of $\mathbf{J} \times \mathbf{B}$ needs these corrections to accurately model hydromagnetic interactions. We have also seen that the “compatible energy” method still produce improvements in most problems, but that the separate evolution of \mathbf{v} and \mathbf{B} creates an ambiguity in the allocation of kinetic, thermal, and magnetic energy. A form of SPMHD that evolves these two variables in a more compatible way would improve this method and could pave the way for a completely compatible treatment of hydromagnetic energy.

In short, SPH and SPMHD produce believable results for simple MHD flows, but in order for them to be more thoroughly studied and tested, the following issues must be addressed:

- The tensile instability is clearly a significant factor in the choice of simulation parameters and must be dealt with in a way that does not destroy the accuracy of the solution.
- The artificial dissipation terms must be limited so that solutions are not oversmoothed at the expense of convergence.
- A better MHD energy conservation scheme should be developed so that the importance of the noisy ∇h terms can be characterized in more detail. In particular: is this correction needed when energy is conserved?

- Since magnetic diffusion is anisotropic in plasmas, an efficient formulation for such anisotropic diffusion is needed. This could come in the form of a more efficient MLPG method (with streamlined numerical quadrature rules) or from some anisotropic SPH discretization.

In any case, we have obtained solutions that are convergent within a range of parameters. In the next chapter we will use this method to model actual laboratory plasmas in the Compact Toroid Injection eXperiment.

Chapter 5

Simulations

In this chapter we simulate the compression of a plasmoid in the U. C. Davis Compact Toroid Injection eXperiment (CTIX) [6] and compare our results to experimental data. In particular, we show that there is a general consistency between the simulations and experimental data.

5.1 The Device

The CTIX device creates and accelerates magnetized plasmas. It has been used to study processes in plasma acceleration [141], magnetic reconnection [25], and fuel injection for magnetic confinement devices [142].

Figure 5.1 shows a schematic of CTIX. The most significant geometrical features on the device are a pair of concentric electrodes through which current flows to create and accelerate the plasmoid. There are four sections of the device:

- the *formation section* where the plasma is created.
- the *acceleration section* in which the plasma is pushed forward by strong electric currents.
- the *drift section* in which the acceleration ceases and the plasma drifts freely to the vessel. A compres-

sion nozzle has been placed in this section to focus the plasma as it stops accelerating; when we study this focusing effect, we refer to this as the *compression section*. In any case, our simulations focus on this section of the device.

- the *target vessel*, which offers a view of the plasma to various diagnostics and to bystanders with sufficiently high metabolic rates.




Figure 5.1: The CTIX device (compression nozzle not shown)

In the formation section, gas is puffed into a chamber, and a capacitor bank C_f is charged until the voltage between the inner and outer electrodes is $V_c \approx 100$ kV. A plasma discharge is formed from the gas from the inner to the outer electrode which ionizes it and creating a plasma. This plasma is allowed to equilibrate over a period of a few microseconds into a state with an internal magnetic field \mathbf{B}_{int} . We discuss this equilibrium state in the next section.

A capacitor bank C_a in the accelerator section then discharges current into the inner electrode along the z axis. This current runs radially outward through the plasma. The current returns along the z axis, completing a circuit with C_a . This current creates a magnetic field \mathbf{B}_p whose interaction with \mathbf{B}_{int} pushes the plasma axially forward. Alternatively, one may think of this acceleration as resulting from $\mathbf{J} \times \mathbf{B}_{int}$, where \mathbf{J} is the density of the current I running through the plasma. Figure 5.2 illustrates the acceleration of the toroidal

plasma confined by the inner and outer electrodes in the accelerator section.




Figure 5.2: The toroidal pushing field, distinguished from the plasma's internal field. The pushing field is created by electric current, which flows through the outer electrode, passes through the conducting plasma, and returns along the inner electrode.

In any case, this acceleration continues for as long as the plasma connects the two electrodes, with I determined by the properties of the circuit consisting of the capacitors, the electrodes, and the plasma. The inner electrode terminates at the beginning of the drift section. Once the bulk of the plasma detaches from the inner electrode, it is no longer accelerated and drifts freely until it reaches the compression section. The compression of the plasmoid in this section is the subject of our investigation.

5.1.1 Diagnostics

CTIX has several apparatuses designed to probe the plasma and magnetic field within the vessel. In this section we will only discuss those diagnostics which can be used to validate our resistive MHD model. Since the simulations we have run for this thesis predict large, qualitative effects, the diagnostics need only provide relatively coarse information.

Magnetic Probes

Magnetic probes [22] gauge the strength of a magnetic field by measuring the voltage induced in a loop of wire immersed in that field. The voltage V is related to the magnetic flux in the loop by Faraday's




Figure 5.3: The compression nozzle with ports for magnetic probes.

Law:

$$V = -\frac{\partial}{\partial t} \left(\int \mathbf{B} \cdot \mathbf{n} dS \right) = -\frac{\partial B_n}{\partial t} A \quad (5.1)$$

where A is the cross-sectional area of the loop in which current is induced, \mathbf{n} is a unit vector normal to the current loop, and $B_n = \mathbf{B} \cdot \mathbf{n}$. The voltage from the probe is filtered and integrated in time to produce the field measurement. Probes of this type are used to measure B_θ and B_z near the boundary of the vessel in the compression nozzle and in the accelerator section. The accelerator probes are shown in Figure 5.1 and are located at $z = 57$ cm, $z = 91$ cm, and $z = 142$ cm. The nozzle contains 9 probes in a 3×3 array of equally-spaced axial and azimuthal positions as shown in Figure 5.3.

The target vessel also has 3 of these probes in recessed cavities which have been used in the original “drift” configuration of the experiment without the nozzle.

The field signals from each probe are acquired and integrated in time, whereupon they can be used to analyze the structure of the plasma as it passes through the device. The magnetic field is localized within the plasma, so the field signals indicate the plasma’s position at a given time. Figure 5.4 shows the progression of a typical CTIX compact toroid with B_z measured at three axial positions and B_θ measured at two (the first

B_θ probe had to be removed to make room for other instrumentation in our experiment. These signals are characteristic of a well-formed plasma.

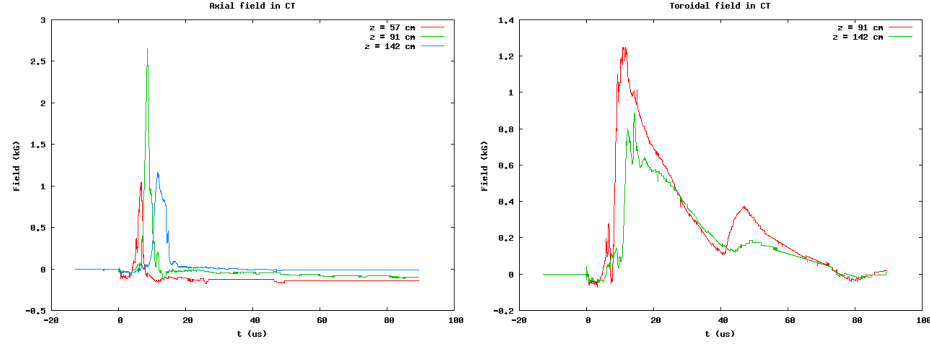


Figure 5.4: (a) B_z and (b) B_θ in a well-formed compact toroid as it passes probes at different axial points in the acceleration section. Taken from CTIX shot 71310.

We use the probes as a gauge of how well the CT has formed, and whether it has kept its shape under acceleration and compression. We also estimate the relative strength of the field within the plasma using calibration factors determined by Howard in his PhD thesis [25]. These factors, which relate the strength of the measured field on the outer electrode to the field of the plasma within the vessel, were determined using a simple magnetostatic approximation of the acceleration section.

High-Speed Camera

A high-speed camera sits outside the target vessel, ready to capture visible light emitted from emerging plasma. We use this camera to determine whether the plasma escapes the compression nozzle at all. By adjusting the sensitivity of the camera to exclude ambient light and by keeping the shutter open for a long time after $t = 0$, we can be confident that any light that emerges from the nozzle is plasma that has made its way through.

5.2 Plasma Compression

Previous experiments with CTIX have demonstrated that the plasmoids it generates retain their magnetic structures while they are being accelerated, but that this structure is quickly dismantled once the plasma passes the center electrode and begins to drift. This is largely because the plasma expands at the Alfvén speed $v_A = B/\sqrt{\mu_0\rho}$ and becomes considerably less dense. The CTIX experiment has a thermal-to-magnetic pressure (β) of a few percent, meaning that the expansion of the plasma is dominated by the presence of the magnetic field. Thus, as the field increases in strength, the expansion rate of the plasma increases rapidly. This expansion makes this method of acceleration unattractive for applications that require the delivery of material to a specific target. In unmagnetized plasmas, the expansion is purely hydrodynamic, but these plasmas cannot be accelerated by magnetic fields.

If the plasma can be successfully focused into a nozzle such as the one shown in Figure 5.3, the density and magnetic field strength can be increased and the plasma may retain some of its structure. As we will see from our simulations and from the associated experimental data, the magnetic field of the plasma can be a significant obstacle to the motion of the plasma. In ideal MHD, the behavior of the magnetic field is essentially governed by magnetic flux conservation: as magnetized material is compressed, the field becomes more intense so that the magnetic flux through a patch of material is conserved. Thus we would expect \mathbf{B} to intensify as the plasma is compressed during its passage through the nozzle. The presence of resistivity allows the field to diffuse through the plasma, which suggests that a cool plasma will compress more easily. However, this same diffusion allows the magnetic field to reconnect with itself, creating new magnetic structures that can also impede the flow of the plasma. Our simulations predict the emergence of such structures in the center of the vessel.

In the simplest terms, the balance of kinetic to magnetic energy in a low β plasma like that of CTIX can determine whether a plasma successfully emerges from the nozzle. After all, we expect that less magnetic field should provide less resistance. However, we will see that this is not the whole picture, and that the specific configuration of the magnetic field within the plasma can significantly affect this process.

5.3 Simulating the Compression of a Plasmoid

The behavior of the plasma in the CTIX device is interesting and largely unknown. The internal magnetic field retained by the CT after acceleration determines its rate of expansion in the drift section and its magnetic pressure in the compression section, but can only be very crudely measured at the outer electrode. In particular, the ratio of kinetic to magnetic field energy in the plasmoid determines how far it progresses through the compression nozzle. The kinematic pressure of the CT is essentially insignificant in the acceleration and drift sections, but it is not clear how much this changes under compression. In short, the accurate description of a plasmoid's journey through CTIX is a complicated topic of ongoing research. It is unlikely that the results of any simulation will accurately describe the dynamics without years of dedicated effort.

The compression of the plasma by the nozzle presents an interesting problem that can be addressed by simulation without focusing on the detailed structure of the plasma. We have mentioned that the plasma's kinetic and magnetic energies determine its progress through the nozzle. As the internal magnetic field of the CT increases, the magnetic pressure $B^2/2\mu_0$ increases rapidly. At some point, the magnetic field prevents the passage of the CT through the nozzle. This threshold depends upon the geometry of the nozzle and the ratio of the kinetic to magnetic energy in the CT. Other factors, such as the temperature of the plasma, which governs the diffusion of the magnetic field, also factor into the threshold. In this section we describe the results of a series of simulations involving plasmoids with various axial velocities and magnetic fields in pursuit of a better understanding of the compression process.

In our simulations, we represent a CT with a cluster of nodes, and we do not represent the vacuum explicitly. This is justified to the extent that the magnetic field is zero outside the plasmoid. Experimental data suggest that this is the case. This allows us to focus all of our computational effort on resolving the plasma. The Lagrangian formulation of SPH is well suited to this problem.

5.3.1 Initial Conditions: Force-Free Equilibrium

The plasmas generated by CTIX are compact toroids (CTs), which are donut-shaped blobs with self-contained magnetic fields that are localized in space. These built-in magnetic fields make CTs easy to manipulate with external currents. The structure of the magnetic field in a CT is probably very complicated in practice, but can be understood in terms of a relatively simple force-free magnetostatic equilibrium of the plasma between the concentric conducting electrodes of the CTIX device. A force-free magnetic configuration is any for which $\mathbf{J} \times \mathbf{B} = \mathbf{0}$, or

$$\nabla \times \mathbf{B} = \lambda \mathbf{B} \quad (5.2)$$

In his PhD thesis, Howard [25] gives an excellent overview of a cylindrical force-free configuration for CTIX plasmoids, which is referred to in the literature as the Taylor state. The components of a toroidal Taylor-state magnetic field in a compact toroid located between z_1 and z_2 with inner and outer radii r_1 and r_2 are, in cylindrical coordinates:

$$\begin{aligned} B_r(r, z) &= B_0 \frac{k_z}{\lambda} \cos[k_z(z - z_1)][J_1(k_r r) + fY_1(k_r r)] \\ B_\theta(r, z) &= -B_0 \frac{k_r}{\lambda} \sin[k_z(z - z_1)][J_0(k_r r) + fY_0(k_r r)] \\ B_z(r, z) &= -B_0 \sin[k_z(z - z_1)][J_1(k_r r) + fY_1(k_r r)] \end{aligned} \quad (5.3)$$

where $J_{0,1}$ and $Y_{0,1}$ are Bessel functions of the first and second kind and $k_r^2 + k_z^2 \equiv \lambda^2$ are the separation constants related to the magnetic eigenvalue λ in (5.2). The values of k_r and f are determined by the condition that $\mathbf{B} \cdot \mathbf{n} = 0$ at the inner and outer electrodes, i.e. $B_r(r_1, z) = B_r(r_2, z) = 0$. The radial field vanishes at the electrodes when $f = -J_1(k_r r_1)/Y_1(k_r r_1) = -J_1(k_r r_2)/Y_1(k_r r_2)$. Thus, to find k_r , we solve the nonlinear equation

$$J_1(k_r r_1)Y_1(k_r r_2) - J_1(k_r r_2)Y_1(k_r r_1) = 0. \quad (5.4)$$

The CTIX device has $r_1 = 0.0334$ m and $r_2 = 0.0762$ m, so $k_r = 75.18 \text{ m}^{-1}$. k_z is determined by the requirement that $\mathbf{B} \cdot \mathbf{n} = 0$ at the axial ends z_1 and z_2 of the CT, or that $B_z(r, z_1) = B_z(r, z_2) = 0$, so $k_z = m\pi/(z_2 - z_1)$. The observed CTs are about 1 m long, which corresponds to $k_z = 3 \text{ m}^{-1}$.

The Taylor state is only a force-free configuration when the CT is confined within conducting walls at r_1 , r_2 , z_1 , and z_2 . In this sense, (5.3) is only an approximation of the initial conditions for a CTIX plasmoid. For now, this approximation remains our best guess. Other idealized configurations are possible: for example, a reverse field configuration (RFC) is essentially a CT without the toroidal field component B_θ and is the focus of some other plasma acceleration efforts.

We initialize a CT as a toroid with inner and outer radii r_1 and r_2 spanning the axis along $[z_1, z_2]$. For simplicity, the toroid has a uniform density ρ_0 , an axial velocity $\mathbf{v}_0 = (0, 0, v_{z0})$, a uniform specific thermal energy u_0 , and a magnetic field specified by the Taylor state with k_r , k_z determined by the geometry of CTIX and B_0 characterizing the strength of the magnetic field.

The typical speed of a CTIX plasmoid is around 150 km/s, and the magnetic fields generated within measure between 2000 - 12000 G at the outer electrode. The plasmas vary greatly in density, but do not exceed number densities of $10^{15}/\text{cc}$. Howard estimates that typical ion and electron temperatures in the plasma fall into the range of $k_B T = 10 - 100 \text{ eV}$ [25], where $k_B T$ is a temperature as measured in electron Volts and k_B is the Boltzmann constant. We typically initialize the thermal energy in the plasma by computing SPH nodal masses from the total mass of the plasma and then finding the specific thermal energy $u = k_B T / m$ for each node.

For the simulations in this chapter, we have assumed azimuthal symmetry. This is an idealization, but should hold mathematically for initial configurations with no azimuthal irregularities. To enforce this symmetry, we compute solutions to the resistive MHD equations in the first quadrant $0 \leq \theta \leq \pi$ using the azimuthal symmetry boundary condition we described in previous chapters. The actual dimensions of the compression nozzle are used in our simulation. The length of the nozzle is 45 cm, and the large and small ends have radii of 6.391 and 2.013 cm respectively. We initialize CTs of length 5 cm. This is much shorter than the length of a CT as observed during acceleration, but we are interested in the behavior of an idealized

plasmoid as it enters the nozzle. We have run our simulations at relatively low resolutions, using initially-uniformly spaced nodes numbering between 20 and 40 along the axis of a CT.

5.3.2 The Plasma/Vacuum Interface

In Chapter 4 we described an implicit magnetic diffusion algorithm that produced accurate and stable results provided that boundary conditions were appropriately specified for \mathbf{B} . In this situation, we do not know the value of \mathbf{B} at the boundary—we only know that $\nabla \times \mathbf{B} = \mathbf{0}$ at the plasma-vacuum interface. We have seen that the SPH magnetic diffusion equation satisfies $\nabla \times \mathbf{B} = \mathbf{0}$ automatically on free boundaries, but the implicit magnetic diffusion method requires us to formulate constraint equations involving the values of \mathbf{B} near the boundary. However, we have found that specifying its curl here using (3.25) or (3.28) can generate ill-conditioned matrices. To avoid this problem, we use the completely explicit magnetic diffusion variant given by (4.27) for the simulations in this chapter. This requires a greater constraint on the time step to guarantee the stability of \mathbf{B} . Fortunately, we have found that the time steps determined by the artificial dissipation models for fast-moving magnetized plasmas are sufficiently small to produce good results for the simulations in this chapter. This may not be the case at much higher resolutions, though, so this remains an area of concern.

The ideal MHD contribution to $\frac{D\mathbf{B}}{Dt}$ is numerically evaluated using the gradients involving sums over neighbor contributions. These gradients are very inaccurate on boundaries, where these contributions are truncated. $\frac{D\mathbf{y}}{Dt}$ is affected in the same way near free boundaries. If one ignores these inaccuracies, magnetic and kinetic energy can build within the plasma at the free boundary. This energy is a pure numerical artifact and can produce highly magnetized plasmas even in cases where the magnetic field is extremely weak.

The difficulty of modeling the interface between a plasma and a vacuum is well-established. Numerically the vacuum is a very stiff region of the problem since the magnetic field diffuses through it at the speed of light. In this light, most authors choose to model the vacuum as a very poor conductor, and proceed to solve the magnetic diffusion equation everywhere in space. The numerical growth in the field and the energy comes from the fact that we simply do not model the vacuum and its associated magnetic field. We

have found that this growth can be quelled by setting $\frac{D\mathbf{B}}{Dt}$ to zero on the nodes that compose the free boundary. Since the field is frozen into the material in the ideal MHD approximation and since these free boundaries are not under compression, this prescription makes some sense. Zeroing the portion of $\frac{D\mathbf{v}}{Dt}$ from $\mathbf{J} \times \mathbf{B}$ is also justified on these nodes, since the current density there is zero by definition.

5.3.3 Plasma Resistivity

The simplest model of the electrical resistivity of a plasma is the Spitzer model for hydrogen, which depends heavily on the electron temperature of the plasma [70]:

$$\eta = 5.2 \times 10^{-15} \frac{Z \ln \Lambda}{T_e^{3/2}} \quad (5.5)$$

where η is measured in $\Omega \text{ m}$ and T_e is in eV. Spitzer derived this model by expressing the resistivity in terms of the electron-ion collision frequency and computing that frequency from simple Coloumb scattering. The marked sensitivity of η to T_e means that the resistivity of a plasma can vary dramatically in space and time. So that we may study the simplest dynamics of the plasma, we perform our studies with constant resistivities that we compute by using a prescribed T_e in (5.5).

5.3.4 Artificial Dissipation and Spurious Viscous Heating

Preliminary simulations of the CT have shown that the relatively high axial velocity of the CT produces significant artificial dissipation. This is not surprising, since our test problems have already indicated that the existing machinery for artificial dissipation in SPH is crude and needs refinement. In any case, there will always be a gradient in the velocity near the surface of the plasma since $\mathbf{v} = \mathbf{0}$ in the vacuum. The dissipated kinetic energy is converted to thermal energy as usual, but the thermal energy produced from the large dissipation rapidly superheats the simulated plasma. This causes the kinematic pressure to skyrocket, and the computation quickly deviates from observed experiment. We emphasize here that this heating is purely numerical—it is not physical and is not related to Joule heating in the plasma.


We have stated in previous chapters that the standard artificial dissipation models in SPH are very crude and tend to convert too much kinetic energy into thermal energy. This superheating is a clear demonstration of the crudeness of these numerical dissipation schemes. An item of great interest in the SPH community is the development of a limiter that curbs such excess dissipation while sufficiently smoothing discontinuities. Such a limiter would reduce the numerical superheating we have observed in these simulations.

In reality, plasmas at high temperatures shed excess thermal energy by radiation. Howard's thesis shows several pictures of the radiation from CTs in this experiment. The simplest way to prevent the artificial superheating is to drop the dissipated kinetic energy instead of converting it, with the understanding that the energy lost has been radiated away. A more physically legitimate method is to treat the plasma as a black body and use the Stefan-Boltzmann Law to determine the amount of heat lost to radiation. We do not use radiative cooling in our simulations because we cannot be certain that the equilibrium-based assumptions of resistive MHD model the plasma's temperature reliably. The heat added by artificial dissipation only compounds this uncertainty. Thus, we discard the artificial heat as mentioned above. This causes energy loss that is qualitatively similar to radiative heating but that serves mainly to curb this numerical problem without introducing new physics.

5.3.5 Results

We are interested primarily in whether a plasmoid passes through the nozzle successfully. We initialize a toroid-shaped plasma of nodes that are initially uniformly spaced on a hexahedral grid with a uniform mass density of $\rho_0 = 1.7 \times 10^{-6} \text{ kg/m}^3$. The initial velocity of the plasma is 150 km/s, and we assign this velocity to each SPH node. Its initial temperature is 20 eV, which corresponds to an initial specific thermal energy of $2.9 \times 10^9 \text{ J/kg}$. Figure 5.5 shows a volume rendering of the plasma's initial configuration, with the colors representing the mass density and the arrows the large axial fluid velocity.

This figure and the others in this chapter were produced by interpolating the SPH nodes in the plasma to a stationary tetrahedral mesh that fills the volume of the nozzle. The contours show surfaces of equal values. The simulations were conducted over a time interval of $3\mu\text{s}$, which corresponds roughly to the




(a)

Figure 5.5: The initial configuration of the plasma as rendered by the Visit visualization tool. The nozzle boundary is represented by a purely cosmetic uniform structured mesh.


experimentally-observed transit time of the plasma light in the experiment. The visualizations were generated after every 50 ns of simulation time. We use the excellent Visit 3D visualization tool [143] developed at the Lawrence Livermore National Laboratory to generate many of the images in this chapter.

We have run several simulations with varying magnetic field strengths, and the results suggest that even a modest magnetic field can prevent a CT from completing its journey. Figure 5.6 shows the mechanism by which a 5000 G magnetic field inhibits the passage of the plasmoid. The plasma expands rapidly in the radial direction very early on, and its magnetic field is allowed to reconnect with that of the plasma expanding in the opposite direction. This reconnection causes the magnetic field to migrate from the plasma to the center of the nozzle, where it proceeds to push the material outward. The plasma collects along the edge of the nozzle. This makes it very difficult for it to escape. It also effectively evacuates the material from the center of the vessel. We will see in the next section that the experimental data corroborates this prediction.

We see magnetic reconnection even in simulations for which we do not include physical resistivity. This is because we use artificial resistivity to damp oscillation in \mathbf{B} , as described in Chapter 4. This diffusion would still occur in some capacity even if the simulations were run at much higher resolution, since our results from the test problems in the last chapter indicate that the artificial resistivity models we use exhibit



(a)



(b)

Figure 5.6: The development of a large magnetic field along the axis pushes the plasma outward along the nozzle. (a) shows the magnetic field and (b) the plasma's mass density.

excessive diffusion. Because of this, it is unlikely that we can use our method in its current form to model truly ideal magnetohydrodynamics. For modeling the cold plasmas in CTIX, this does not matter, since significant magnetic diffusion is likely to be present at all times.

Nor does lowering the plasma temperature appear to change the behavior of the magnetic field. We have used the CTIX device to form several compact toroids and accelerated them into the nozzle, and the data we have obtained from these experiments shows that these plasmas do not successfully complete their journey. Figure 5.7 shows photos of the nozzle taken by the camera after the plasma enters the nozzle. We have reduced the sensitivity of the camera in order to reduce any ambient light from plasma inside the nozzle. To ensure that all escaping plasma was captured, we left the shutter on the camera open for 1 ms, which is over an order of magnitude longer than the travel time for the plasma. None of the plasmas that we formed escaped from the nozzle.

We have estimated the magnetic field components of the CTs, which are shown in Table 5.1. The field strength of the plasma, shown in the table, is related to the magnetic probe measurement by

$$\mathbf{B}_{plasma} = C\mathbf{B}_{probe}, \quad (5.6)$$

where $C = 2.32853$ is the calibration factor for the last probe in the acceleration section.

Shot	B_z (T)	B_θ (T)
71310	0.27	0.20
71336	0.13	0.14
71337	0.068	0.17
71361	0.24	0.27

Table 5.1: Estimated peak magnetic field measurements for shots in Figure 5.7. These fields are measured as the plasma enters the nozzle.

If a much weaker magnetic field is used, the plasmoid can pass all the way through the nozzle. Figure 5.9 shows the passage of a compact toroid with an internal magnetic field of around 2000 G. The total transit time of the CT is around $3 \mu\text{s}$. The plasma's passage is impeded greatly by any increase in the magnetic field. Unfortunately, the plasma acceleration process in CTIX seems to require a magnetic field in the plasma that exceeds the threshold for successful compression in this geometry.






Figure 5.7: Experimental photographs of the nozzle after the plasma compressions. Shots (a) 71310, (b) 71336, (c) 71337, and (d) 71361.




(a)




(b)

Figure 5.8: The final configuration of a plasma with $B_0 = 5000$ G that does not escape from the nozzle. (a) mass density, and (b) magnetic field intensity.



(a)



(b)

Figure 5.9: The emergence of a plasma with $B_0 = 2000$ G from the nozzle. The magnetic field growth along the axis is still present, but the kinetic energy of the plasma allows it to overcome the magnetic pressure. (a) mass density, and (b) magnetic field intensity.

We have run several low-resolution simulations of these CTs in an attempt to find a correspondence between the balance of kinetic energy to magnetic energy required for a successful traversal of the nozzle. The results of these simulations, characterized by the velocity and the internal magnetic field of the compact toroids, are summarized in Table 5.2.

v_z (km/s)	B_0 (T)	pass	v_z (km/s)	B_0 (T)	pass	v_z (km/s)	B_0 (T)	pass
100	0.1	yes	150	0.1	yes	200	0.1	yes
100	0.2	no	150	0.2	yes	200	0.2	yes
100	0.5	no	150	0.5	no	200	0.5	no
100	1.0	no	150	1.0	no	200	1.0	no

Table 5.2: Results of simulations with various axial velocities and internal magnetic field strengths. The “pass” column indicates whether the plasma passes successfully through the nozzle.

Magnetic tension and reconnection

The simulations have predicted the emergence of a strongly magnetized mass in the center of the vessel. To study this process more closely, we have taken a cross section of the CT and run higher-resolution 2D simulations of the magnetic reconnection process. The plasma expands radially at the Alfvén speed, carrying the magnetic field with it. The field reconnects with itself in the center and then grows more and more intense until it is able to repel any further inflowing plasma. If the initial field is strong enough, the center region can be strongly magnetized without containing much plasma. Thus, the CT is pushed against the sides of the nozzle, where it sacrifices kinetic energy for thermal energy.

The reconnection of the field depends mainly on the axial component of \mathbf{B} at the inner radius of the plasmoid. In a CT or a reverse field configuration (RFC), the axial component of the magnetic field is itself axially symmetric. The axial and radial components compose the *poloidal* magnetic field. As the plasma flows to the center of the vessel, the axial magnetic field vectors are aligned, and as they reconnect they produce a stronger field. Figure 5.10 shows how the poloidal magnetic field lines can reconnect in these circumstances. A 2D simulation describing the development of the field in the expanding CT is pictured in Figure 5.11.

Compare this process with a plasma jet in which there is a purely *toroidal* magnetic field $\mathbf{B} =$





Figure 5.10: Schematic representation of *poloidal* magnetic reconnection in a compact toroid that is expanding radially and moving to the right. The poloidal (axial) field lines are aligned as they approach one another, reconnect, and produce a more intense field.



(a)

(b)

Figure 5.11: Magnetic field amplitude in 2D simulation of a cross section of an expanding CT. (a) initial configuration, and (b) after reconnection ($t = 0.1\mu\text{s}$).




Figure 5.12: Schematic representation of toroidal magnetic reconnection. The toroidal field lines are opposed as they approach one another, annihilating one another upon reconnection.

$(0, B_\theta, 0)$ in which the field vectors are opposed to one another across the axis. As the plasma flows inward, these magnetic field vectors annihilate one another and no central field is produced. Figure 5.12 illustrates the reconnection of these magnetic field lines. 2D cross-section simulations of such jet configurations confirm this intuitive picture: when the magnetic field has no significant axial component, no significant magnetized region is created in the center of the vessel. Figure 5.13 shows the results of such a 2D planar simulation of the radial inflow of a toroidal jet.

Within this magnetic reconnection and growth, the plasma is not pushed with such force to the sides of the nozzle. It appears that a plasma jet should have little to no difficulty traversing the nozzle, regardless




Figure 5.13: Magnetic field amplitude in 2D simulation of a cross section of an expanding jet. (a) initial configuration, and (b) after reconnection ($t = 0.1 \mu\text{s}$).

of magnetic field strength. Our 3D simulations demonstrate that this is the case.

Experimental evidence confirms that jets can pass through the nozzle relatively unimpeded. The CTIX device can be configured to produce a plasma jet simply by disabling the solenoid that produces the axial field. We confirmed the jet field configuration by inspecting the magnetic probe signals. The progression of a typical jet through the accelerator region is shown in Figure 5.14. The axial magnetic field of the jet is much smaller is oriented opposite to that of a compact toroid.




Figure 5.14: (a) B_z and (b) B_θ in a toroidal plasma jet as it passes probes at different axial points in the acceleration section. This is experimental data taken from CTIX shot 71356. Contrast these signals with those in Figure 5.4.




Figure 5.15: Experimental photographs of plasma jets emerging from the nozzle at right. Shots (a) 71356, (b) 71359, (c) 71360, and (d) 71375.

Several shots were taken in this configuration at varying field strengths, and the images from the camera indicate that plenty of plasma escapes from the nozzle in each case. Figure 5.15 shows the emergence of strongly magnetized jets from the nozzle. The estimated magnetic field strengths are shown in Table 5.3. Evidently, even jets with very strong toroidal magnetic fields have no difficulty emerging from the nozzle. This behavior can be understood in some sense by the notion of magnetic tension. The axial magnetic field and radial magnetic field components define the “poloidal” magnetic field, which is responsible for holding plasmoids together. Without poloidal field, a magnetized plasma simply blows itself apart. The magnetic tension $\mathbf{B} \cdot \nabla \mathbf{B}$ associated with the poloidal field is thus both a help and a hindrance to a plasmoid: just as

Shot	B_z (T)	B_θ (T)
71356	0.0018	0.32
71359	0.018	0.53
71360	0.053	0.15
71375	0.15	0.32

Table 5.3: Estimated peak magnetic field measurements for shots in Figure 5.15. These fields are measured as the plasma enters the nozzle.

magnetic tension preserves the shape of the CT and allows plasma and magnetic field to move together, the same tension prevents the plasma from assuming arbitrary shapes.

5.4 Conclusion

We have shown that CTs and other plasmoids with axial (“poloidal”) magnetic fields have great difficulty getting through a nozzle without acceleration. This difficulty is caused by the inward expansion of the plasma and the resulting reconnection of the magnetic field at the center of the vessel. One way to improve the compression process would be to have a center electrode in the compression section. This would both continue the acceleration of the plasma and prevent it from reconnecting with itself.

Currently, we are not able to satisfactorily simulate the acceleration of plasma with SPH, so such an experiment falls outside the scope of this thesis. In particular, the magnetic field in the vacuum regions would have to be modeled. It is possible to do this with stationary nodes of extremely little mass. The MLPG method also offers a way of solving the magnetostatic equation for the vacuum field with great accuracy near conducting boundaries. Such a method could be coupled to the SPH plasma model to find a global magnetic field and represent an equivalent circuit that delivers energy to the plasma from the CTIX capacitor bank.

In this chapter, we have applied our SPH-based resistive MHD method to the compression of a plasma by a nozzle and obtained results that are consistent with experimental data. These results have been published in a peer-reviewed journal [144]. While there is much left to be studied in both the numerics and the plasma dynamics studied in this chapter, we believe we have shown that the numerical method is capable of producing plausible answers to difficult problems in laboratory plasma physics.

Chapter 6

Conclusion

In this chapter we summarize the method we have described: what it does well, and what it does not. We then discuss the most promising directions of future development.

6.1 A Critical Assessment of SPMHD

This thesis is a both a quantitative evaluation of existing methods in treating magnetic fields with SPMHD and an extension of the method to simulate laboratory plasmas. The curious reader will not find convergence studies like those we have done published anywhere previously in the SPMHD literature. In studying the convergence of these methods, we have strived to illustrate both their strengths and shortcomings, and to invite interested practitioners to improve upon them where it is most necessary to do so. We have also provided software that can be used to reproduce all of the results—convergence studies and all—with relative ease so that they may be seen firsthand. This software is described in Appendix B.

There seems to be a general reluctance in the SPH community to subject these methods to rigorous convergence testing. Very few authors have made it a regular practice in their work. In essence, the rate of convergence of a numeric method is *the* objective measure of that method’s reliability: it is literally the rate at which the solution improves with more computational effort. The lack of convergence studies in the current

literature certainly has harmful effects: not only does it foster skepticism about SPH in other camps, but it also allows the continuation of dubious practices, such as the routine use of values of n_h that are too low to achieve convergent results. We hope that the presence of these numbers and the ability to reproduce them will encourage others in the SPH community to adopt the use of convergence studies as a method for evaluating algorithmic improvements in the future.

The most significant result of our study of SPH is that the method does not produce convergent results to many simple hydrodynamic test problems at values of n_h that are typically used to run simulations in the field. Alarming, multidimensional hydrostatic equilibrium cannot be maintained when n_h is 1.2 or 1.4. The propagation of compressional waves approaches first-order when $n_h = 2$. This indicates that their resolution is limited by artificial dissipation and possible node disorder. However, at lower values of n_h the convergence is destroyed by large phase errors. Finally, shocks are much better resolved when $n_h = 2$, and SPH can achieve first-order accuracy even in the presence of strong shocks. For pure hydrodynamic problems, the $n_h = 2$ appears to yield the most convergent results to our test problems.

In our MHD simulations, things are more complicated. While higher values of n_h tend to increase the convergence rates of the errors in the computed solutions, this convergence is visibly degraded by oscillations probably induced by the tensile instability. This instability appears to be exacerbated by higher values of n_h , which indicates to us that the SPMHD discretizations of the induction equation and the Lorentz force may have significant problems. The competing effects of the non-convergence of low values of n_h and the onset of the tensile instability make it hard to predict what value of n_h is best to use for a particular MHD problem. This is reflected in the various values of n_h used by Price [84], with and without explanation.

Finally, we see that Brookshaw's method for magnetic diffusion produces results that are superlinearly convergent for simple configurations, even when n_h is lower than 2 and when the nodes are disordered. However, its representation of spatially-inhomogeneous resistivities is problematic, and since other authors that use the discretization ([4], [46]) do not rigorously show the convergence of the method, we cannot say whether this problem is limited to its use in magnetic diffusion. The MLPG method handles such dynamic resistivities much more gracefully and can properly treat anisotropic resistivity. However, its computational

performance is not competitive with SPH on moving sets of points. If the numerical quadrature in the MLPG integrals could be accelerated, the MLPG method would be a very viable candidate for modeling magnetic diffusion.

6.2 Extensions

In this thesis we have extended Smoothed Particle Magnetohydrodynamics to accomodate the physics of laboratory plasmas. We explored two mesh-free techniques for solving the magnetic diffusion equation and demonstrated that they are convergent. The SPH formulation has been shown to perform well enough to be used in magnetohydrodynamic flows, and can treat spatially- and temporally-varying resistivities, which is crucial for modeling plasmas with nonuniform temperatures and densities.

We have explored the formulation of boundary conditions near curved surfaces, a topic which has not received much treatment in the SPH literature. Boundary conditions on \mathbf{B} near surfaces and free boundaries have been given and shown to produce believable results. We have also looked at boundary conditions for expressing symmetries in systems in order to improve computational performance.

Finally, we have introduced a more accurate and conservative method for evolving the thermal energy in magnetohydrodynamics problems based on Owen's compatible energy method [112].

These refinements have been used to model the dynamics of magnetized plasma in CTIX [6], an actual plasma experiment, and results have been compared with real data. The compression of a magnetized plasma has been studied and has uncovered one of the great strengths of the meshless fluid approach: that the inherently adaptive resolution of SPH prevents us from having to model the entire geometry of the device.

6.3 Further Study of CTIX

We have begun to study the dynamics of plasmas in the drift section of CTIX, but many questions remain. For example, some experiments seem to indicate that the toroid can tumble off-axis, rendering the system completely non-symmetric. In addition, we have only explored a tiny patch of the parameter space

available to the experiment, both under the circumstances of drift and compression.

The method can also be used to perform rudimentary studies of the acceleration section. It may be possible to mock up the “pushing field” as a source magnetic field on a set of static SPH nodes that interacts with the plasma nodes. The source field could be specified as a function of space and time, or could be consistent with a circuit model that represents the CTIX device.

More simulations like these can be devised as diagnostics improve, allowing the algorithm to be calibrated against the experiment.

6.4 Directions for Improvement

The use of mesh-free methods to solve the magnetohydrodynamic equations is a very recent endeavor undertaken by a very small community. SPMHD is being used by a few in the astrophysics community, but clearly needs some major progress before it is comparable to its grid-based counterparts. Most groups doing simulations of magnetized plasmas in the laboratory use one of a small set of workhorse codes.

There are several options for those looking to improve this method. The artificial dissipation terms we discussed in Chapter 3 and Chapter 4 are clearly too heavy-handed and need to be tempered. Of great help would be a limiter in the form of a variable switch that accurately decides which regions are linear and which are nonlinear, smoothing only the latter and leaving the former unaltered. In our work, we have briefly explored the use of gradient limiting to try to determine regions that are linear and found that even our rudimentary models show noticable improvements in convergence in hydrodynamic and magnetohydrodynamic shock tube problems. A more formal higher-order upwinding scheme would see great improvements in the accuracy of the method.

The engineering community has contributed many improvements to the SPH method itself, some of which solve the consistency problems we discussed in Chapter 3. Perhaps the most promising among these is Corrected SPH, which is gaining ground ([86], [131]). This method is much more accurate than “regular” SPH near boundaries, and goes some length toward addressing the tensile instability [80] that causes so much

trouble in magnetized flows. This method should be extended to solve compressible fluid flows, after which the treatment of curved boundaries and boundary conditions for \mathbf{B} could be further improved.

For the treatment of complex boundaries, the Multiple Boundary Tangent method as described by Yildiz, Rook, and Suleman [120] shows some promise. We have run some simple calculations with a preliminary implementation of this method. However, its approach to representing boundary conditions is a significant departure from the other methods we have used, and its further study (particularly its application to three-dimensional compressible flows) is beyond the scope of this thesis. There are several problems in laboratory plasma physics that could benefit from the ability to define complex boundaries with discrete sets of points and normal vectors. For example, the magnetic probes in the CTIX device are housed in wells within the conducting electrodes and the nozzle. The degree to which the geometry of these wells affects the measurement of \mathbf{B} is an interesting and relevant problem.

Accuracy and stability are not the only limitations of the present method. It seems to be very difficult to represent \mathbf{B} outside of material when the field itself is only defined on material points. While our meshless treatments of magnetic diffusion can deal with resistivities spanning several orders of magnitude (arguably including regions of vacuum), the results we have seen indicate that there is more work to be done on vacuum fields.

A drastic departure from this one would be to represent \mathbf{B} on a grid or a stationary set of points while retaining some form of SPH to solve the fluid equations. This hybrid approach was the idea behind our initial efforts at modeling terrestrial plasmas and simplifies the treatment of magnetic fields near boundaries as well as allowing the use of time-proven techniques for treating the $\nabla \cdot \mathbf{B} = 0$ constraint. Conceptually, the strength of this technique is evident, since a fluid may be expressed in its “natural” Lagrangian coordinates (2.45) - (2.46) while the magnetic field retains its flux-conserving Eulerian formulation (2.34). Perhaps most significantly, magnetic fields in vacuum regions could be easily represented.

Unfortunately, such a hybrid approach introduces practical difficulties present in both Lagrangian and Eulerian methods: for example, the advection of a magnetic pulse becomes a non-trivial problem, and flux limiting must be used to treat the $\nabla \times (\mathbf{v} \times \mathbf{B})$ term in the induction equation. In addition, ρ and \mathbf{v} must

be interpolated carefully from the fluid nodes to the grid for the solution of \mathbf{B} , after which the magnetic force must be interpolated back to these nodes for acceleration. While this sort of interpolation has been explored in great detail by Particle-In-Cell methods [145], it is not clear that the same rules apply, since the hybrid method maintains \mathbf{B} on the grid from step to step. All of these difficulties are surmountable, but the unknowns make this approach a major undertaking.

In any case, there remains much interesting work to be done on meshless approaches to magneto-hydrodynamics.

Bibliography

1. G. J. Phillips and J. J. Monaghan. A numerical method for three-dimensional simulations of collapsing, isothermal, magnetic gas clouds. *Mon. Not. R. Astron. Soc.*, 216:883–895, 1985.
2. S. Børve . *A numerical study of ionosphere and magnetospheric phenomena using particle simulation techniques*. PhD thesis, University of Oslo, 2001.
3. D. J. Price and J. J. Monaghan. Smoothed particle magnetohydrodynamics iii. multidimensional tests and the $\text{div } \mathbf{b} = 0$ constraint. *Mon. Not. R. Astron. Soc.*, 364:384–406, 2005.
4. L. Brookshaw. A method of calculating radiative heat diffusion in particle simulations. In *Proceedings of the Astronomical Society of Australia*, volume 6, pages 207–210, 1985.
5. S. N. Atluri, H. G. Kim, and J.Y Cho. A critical assessment of the truly meshless local petrov-galerkin (mlpg) and local boundary integral equation (lbie) methods. *Comp. Mech.*, 24:348–372, 1999.
6. D. Q. Hwang, R. D. Horton, S. Howard, R. W. Evans, and S. E. Brockington. Advances in ctix accelerator study. *J. Fusion Energy*, 1:81–84, 2006.
7. I. Newton. *Philosophiae Naturalis Principia Mathematica*. 1687.
8. J. C. Maxwell. *A Treatise on Electricity and Magnetism, Vol. 1*. Dover Publications, 1954.
9. P. A. Davidson. *An Introduction to Magnetohydrodynamics*. Cambridge University Press, 2001.
10. G. Valyavin, T. Burlakova, S. Fabrika, and D. Monin. Magnetic fields of white dwarfs. *Astron. Rep.*, 47:587–599, 2003.
11. D. J. Price and S. Rosswog. Smoothed particle magnetohydrodynamics i. algorithm and tests in one dimension. *Science Magazine*, 312:5774, 2006.
12. B. von Rekowski and A. Brandenburg. Outflows and accretion in a star-disc system with magnetosphere and disc dynamo. *Astron. and Astrophys.*, 420:17, 2004.
13. W. H. Bostick. What laboratory-produced plasma structures can contribute to the understanding of cosmic structures both large and small. *IEEE Trans. Plasma Sci.*, 14:703–717, 1986.
14. D. Martinez, C. Plechaty, and R. Presura. Magnetic fields for the laboratory simulation of astrophysical objects. *Astrophys. and Space Sci.*, 307:109–114, 2007.
15. K. Birkeland. *The Norwegian Aurora Polaris Expedition 1902-1903*. Publ. H. Aschehoug and Co., 1908.
16. W. L. Laurence. Physicist creates universe in a test tube, Dec. 12 1956.

17. B. A. Trubnikov. A new hypothesis of cosmic ray generation in plasma pinches. *IEEE Trans. Plasma Sci.*, 20:898–904, 1992.
18. H. Alfven. On the origin of cosmic magnetic fields. *Astrophys. J. Supp.*, 133:1043, 1960.
19. <http://www.iter.org>.
20. K. Moffatt. The amplification of a weak applied magnetic field by turbulence in fluids of moderate conductivity. *J. Fl. Mech.*, 11:625–635, 1961.
21. F. F. Chen. *Introduction to Plasma Physics and Controlled Fusion*. Plenum Publishing, 1984.
22. I. H. Hutchinson. *Principles of Plasma Diagnostics*. Cambridge University Press, 2005.
23. W. H. Bostick. Experimental study of ionized matter projected across a magnetic field. *Phys. Rev.*, 104:292–299, 1956.
24. W. H. Bostick. Experimental study of plasmoids. In *Electromagnetic Phenomena in Cosmical Physics*, volume 6, 1958.
25. S. Howard. *Interaction Dynamics of High Reynolds Number Magnetized Plasma Flow on the CTIX Plasma Accelerator*. PhD thesis, University of California, Davis, 2006.
26. G. Tóth. The $\text{div } \mathbf{b} = 0$ constraint in shock-capturing magnetohydrodynamics codes. *J. Comput. Phys.*, 161:605–652, 2000.
27. S. K. Godunov. A difference scheme for numerical solution of discontinuous solution of hydrodynamic equations. *Math. Sbornik*, 47:271–306, 1959.
28. T. Stix. *Plasma Waves*. Springer Verlag, New York, NY, 1992.
29. D. G. Swanson. *Plasma Waves*. CRC Press, 2003.
30. F. H. Harlow. The particle-in-cell computing method in fluid dynamics. *Meth. Comp. Phys.*, 3:319–343, 1964.
31. R. M. Marder. Gap - a pic-type fluid code. *Math. Comp.*, 29:434–446, 1975.
32. M. Brio and C. C. Wu. An upwind differencing scheme for the equations of ideal magnetohydrodynamics. *J. Comput. Phys.*, 75:400–422, 1988.
33. C. Evans and J. Hawley. Simulation of magnetohydrodynamic flows: A constrained transport method. *Astrophys. J. Supp.*, 332:659–677, 1988.
34. K. S. Yee. Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media. *IEEE Trans. Ant. Prop.*, 14(3):302–307, 1966.
35. A. L. Zachary and P. Colella. A higher-order godunov method for the equations of ideal magnetohydrodynamics. *J. Comput. Phys.*, 99:341–347, 1992.
36. W. Dai and P. R. Woodward. Extension of the piecewise-parabolic method to multidimensional ideal magnetohydrodynamics. *J. Comput. Phys.*, 115:485–514, 1994.
37. J. M. Hyman and M. Shashkov. Mimetic finite difference methods for maxwell’s equations and the equations of magnetic diffusion.
38. R. N. Rieben and D. White. Verification of high-order mixed finite element solution of transient magnetic diffusion problems. *IEEE Trans. Mag.*, 42:25–39, 2006.

39. J. Stone and M. L. Norman. Zeus-2d: A radiation magnetohydrodynamics code for astrophysical flows in two space dimensions: II. the magnetohydrodynamic algorithms and tests. 80:753, 1992.
40. J. M. Stone, T. A. Gardiner, P. Teuben, J. F. Hawley, and J. B. Simon. Athena: A new code for astrophysical mhd. 178:137–177, 2008.
41. C.R. Sovinec, A.H. Glasser, T.A. Gianakon, D.C. Barnes, R.A. Nebel, S.E. Kruger, S.J. Plimpton, A. Tarditi, M.S. Chu, and the NIMROD Team. Nonlinear magnetohydrodynamics with high-order finite elements. *J. Comput. Phys.*, 195:355, 2004.
42. S. E. Kruger, D. D. Schnak, and C. R. Sovinec. Dynamics of the major disruption of a diii-d plasma. *Phys. of Plasmas*, 12:56113, 2005.
43. R. A. Gingold and J. J. Monaghan. Smoothed particle hydrodynamics: theory and application to non-spherical stars. *Mon. Not. R. Astron. Soc.*, 181:375–390, 1977.
44. L. B. Lucy. A numerical approach to the testing of the fission hypothesis. *Astrophys. J. Supp.*, 82:1013–1024, 1977.
45. P. W. Cleary and J. J. Monaghan. Conduction modelling using smoothed particle hydrodynamics. *J. Comput. Phys.*, 148:227–264, 1999.
46. S. C. Whitehouse, M. R. Bate, and J. J. Monaghan. A faster algorithm for smoothed particle hydrodynamics with radiative transfer in the flux-limited diffusion approximation. *Mon. Not. R. Astron. Soc.*, 364(4):1367–1377, 2005.
47. J. P. Morris. *Analysis of smoothed particle hydrodynamics with applications*. PhD thesis, Monash University, 1996.
48. B. Nayroles, G. Touzot, and P. Villon. Generalizing the finite element method: Diffuse approximation and diffuse elements. *Comp. Mech.*, 10:307–318, 1992.
49. C. A. Duarte and J. T. Oden. An h-p adaptive method using clouds. *Comput. Methods Appl. Mech. Engrg.*, 139:237–262, 1996.
50. T. Belytschko, Y. Krogauz, D. Organ, M. Fleming, and P. Krysl. Meshless methods: an overview and recent developments. *Comp. Meth. Appl. Mech. Eng.*, 139:3–47, 1996.
51. S. N. Atluri. *The Meshless Method (MLPG) for Domain and BIE Discretizations*. Tech Science Press, Forsyth, GA, 2004.
52. J. N. Johnson and J. M. Owen. A meshless local petrov-galerkin method for magnetic diffusion. *Comput. and Struct.*, 22:165, 2007.
53. H. Ding, C. Shu, K. S. Yeo, and D. Xu. Development of least square-based two-dimensional finite difference schemes and their application to simulate natural convection in a cavity. *Comp. Fluids.*, 33:137–154, 2004.
54. J. L. Maron and G. G. Howes. Gradient particle magnetohydrodynamics: A lagrangian particle code for astrophysical magnetohydrodynamics. *Astrophys. J. Supp.*, 595:564–572, 2003.
55. O. V. Diyanov. Uncertain grid method for numerical solution of pde. 2008.
56. S. Kitsonias and A. P. Whitworth. Smoothed particle hydrodynamics with particle splitting applied to self-gravitating collapse. *Mon. Not. R. Astron. Soc.*, 330:129, 2002.
57. J. Feldman and J. Bonet. Dynamic refinement and boundary contact forces in sph with applications in fluid flow problems. *Int. J. Numer. Methods Engrg.*, 72:295–324, 2007.

58. P. W. Randles and L. D. Libersky. Boundary conditions for a dual particle method. 83:1476–1486, 2005.
59. D. R. Nicholson. *Introduction to Plasma Theory*. Wiley, New York, NY, 1983.
60. S. I. Braginskii. Transport processes in a plasma. In *Reviews of Plasma Physics*, volume 1, 1965.
61. R. D. Hazeltine and F. L. Waelbroeck. *The Framework of Plasma Physics*. Perseus Books, New York, NY, 1998.
62. J. D. Jackson. *Classical Electrodynamics*. Wiley, New York, NY, 1999.
63. D. Biskamp. *Nonlinear Magnetohydrodynamics*. Cambridge University Press, 1997.
64. G. K. Batchelor. *An Introduction to Fluid Dynamics*. Cambridge University Press, 2000.
65. H. Alfven. Existence of electromagnetic-hydrodynamic waves. *Nature*, 150:405, 1942.
66. Ya. B. Zel'dovich and Yu. P. Raizer. *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena*. McGraw-Hill, London, UK, 1966.
67. D. Ryu and T. W. Jones. Numerical magnetohydrodynamics in astrophysics: Algorithm and tests for one-dimensional flow. *Astrophys. J. Supp.*, 442:228–258, 1995.
68. D. Ryu and T. W. Jones. Numerical magnetohydrodynamics in astrophysics: Algorithms and tests for multidimensional flow. *Astrophys. J. Supp.*, 452:785, 1995.
69. W. Dai and P. R. Woodward. On the divergence-free condition and conservation laws in numerical simulations for supersonic magnetohydrodynamic flows. *Astrophys. J. Supp.*, 494:317, 1998.
70. L. Spitzer and R. Harm. *Phys. Rev.*, 89:977, 1953.
71. P. Coles. *Comments Astrophys.*, 16:45, 1992.
72. J. Stone, J. F. Hawley, C. R. Evans, and M. L. Norman. A test suite for magnetohydrodynamic simulations. *Astrophys. J. Supp.*, 388:415–437, 1992.
73. A. Dedner, F. Kemm, D. Kroner, C. D. Munz, T. Schnitzer, and M. Wessburg. Hyperbolic divergence cleaning for the mhd equations. *J. Comput. Phys.*, 175:645–673, 2002.
74. W. E. Crowley. Free-lagrange methods for compressible hydrodynamics in two space dimensions. In *The Free-Lagrange Method*, volume 238, 1985.
75. W. Benz. 3d models of rotating gas clouds. i - time evolution, mass spectrum and angular momentum. *Astron. and Astrophys.*, 139:378–388, 1984.
76. J. J. Monaghan. Smoothed particle hydrodynamics. *Ann. Rev. of Astron. Astrophys.*, 30:543–574, 1992.
77. G. R. Liu. *Mesh-Free Methods*. 2002.
78. H. J. Weber and G. B. Arfken. *Mathematical Methods for Physicists*. Academic Press, 2005.
79. J. J. Monaghan. Sph without a tensile instability. *J. Comput. Phys.*, 159:290–311, 1999.
80. J. Bonet and S. Kulasegaram. Remarks on tension instability of eulerian and lagrangian corrected smoothed particle hydrodynamics (csph) methods. *Int. J. Numer. Methods Engrg.*, 52:1203–1220, 2001.

81. D. S. Balsara. Von neumann stability analysis of smoothed particle hydrodynamics: suggestions for optimal algorithms. *J. Comput. Phys.*, 121:357–372, 1995.
82. T. J. Martin, F. R. Pierce, and P. A. Thomas. An owner’s guide to smoothed particle hydrodynamics. technical report, 1993.
83. P. W. Randles and L. D. Libersky. Smoothed particle hydrodynamics: Some recent improvements and applications. *Comput. Methods Appl. Mech. Engrg.*, 139:375–408, 1996.
84. D. J. Price. *Magnetic Fields in Astrophysics*. PhD thesis, University of Cambridge, 2004.
85. G. A. Dilts. Moving-least-squares-particle hydrodynamics - i. consistency and stability. *Int. J. Numer. Methods Engrg.*, 44:1115–1155, 1999.
86. J. Bonet and T.-S.L. Lok. Variational and momentum preservation aspects of smooth particle hydrodynamics formulations. *Comput. Methods Appl. Mech. Engrg.*, 180:97–115, 1999.
87. W. K. Liu, S. Jun, and Y. F. Zhang. Reproducing kernel particle methods. *Int. J. Numer. Methods Fluids.*, 20:1081–1106, 1995.
88. F. Colin, R. Egli, and F. Y. Lin. Computing a null divergence velocity field using smoothed particle hydrodynamics. *J. Comput. Phys.*, 217:680–692, 2006.
89. R. J. Thacker, E. R. Tittley, F. R. Pearce, H. M. P. Couchman, and P. A. Thomas. Smoothed particle hydrodynamics in cosmology: A comparative study of implementations. *Mon. Not. R. Astron. Soc.*, 319:619, 2000.
90. D. J. Price and J. J. Monaghan. Smoothed particle magnetohydrodynamics ii. variational principles and variable smoothing length terms. *Mon. Not. R. Astron. Soc.*, 348:128–138, 2004.
91. M. Hamilton and D. Blackstock. *Nonlinear Acoustics*. Academic Press, San Diego, CA, 1998.
92. P. L. Roe. Approximate riemann solver, parameter vectors and difference schemes. *J. Comput. Phys.*, 43:357–372, 1981.
93. R. A. Gingold and J. J. Monaghan. *Mon. Not. R. Astron. Soc.*, 181:375, 1983.
94. J. M. Owen. A tensor artificial viscosity for sph. *JCP*, 201:601–629, 2004.
95. J. P. Morris and J. J. Monaghan. A switch to reduce sph viscosity. *J. Comput. Phys.*, 136:41–50, 1997.
96. J. J. Monaghan. Sph and riemann solvers. *J. Comput. Phys.*, 136:298–307, 1997.
97. E. Chow and J. J. Monaghan. Ultrarelativistic sph. *J. Comput. Phys.*, 134:296–305, 1997.
98. D. J. Price and J. J. Monaghan. Smoothed particle magnetohydrodynamics i. algorithm and tests in one dimension. *Mon. Not. R. Astron. Soc.*, 348:128–138, 2004.
99. D. J. Price. Modelling discontinuities and kelin-helmholtz instabilities in sph. *J. Comput. Phys.*, 227:10040–10057, 2008.
100. O. Agertz, B. Moore, J. Stadel, D. Potter, F. Miniati, J. Read, L. Mayer, A. Gawryszczak, A. Kravtsov, J. Monaghan, A. Nordlund, F. Pearce, V. Quilis, D. Rudd, V. Springel, J. Stone, E. Tasker, R. Teyssier, J. Wadsley, and R. Walder. Fundamental differences between sph and grid methods (preprint). *Mon. Not. R. Astron. Soc.*
101. J. Stoer and R. Bulirsch. *Introduction to Numerical Analysis*. Springer Verlag, New York, NY, 1991.

102. J. J. Monaghan. Implicit sph: Drag and dusty gas dynamics. *J. Comput. Phys.*, 138:801–820, 1997.
103. C. E. Knapp. *An Implicit Smooth Particle Hydrodynamic Code school* =. PhD thesis.
104. R. Courant, K. Fridrichs, and H. Lewy. On the partial differential equations of mathematical physics. *IBM Journal*, pages 215–234, 1967.
105. J. J. Monaghan. On the problem of penetration in particle methods. *J. Comput. Phys.*, 82:1–15, 1989.
106. R. P. Nelson and J. C. B. Papaloizou. Variable smoothing lengths and energy conservation in smoothed particle hydrodynamics. *Mon. Not. R. Astron. Soc.*, 270:1–20, 1994.
107. V. Springel and L. Hernquist. Cosmological sph simulations: the entropy equation. *Mon. Not. R. Astron. Soc.*, 333:649–664, 2002.
108. J. J. Monaghan. Sph compressible turbulence. *Mon. Not. R. Astron. Soc.*, 335:843–852, 2002.
109. J. M. Hyman and M. Shashkov. Natural discretizations for the divergence, gradient, and curl on logically rectangular grids. *Comp. Math. Applic.*, 33:81–104, 1997.
110. G. Rodrigue and D. White. A vector finite element time-domain method for solving maxwell’s equations on unstructured hexahedral grids. *SIAM J. Sci. Comp.*, 23(3):683–706, 2001.
111. E. J. Caramana, D. E. Burton, M. J. Shashkov, and P. P. Whalen. *J. Comput. Phys.*, 146:227, 1998.
112. J. M. Owen. A compatibly differenced total energy conserving form of sph. June 2009.
113. L. Brookshaw. Smooth particle hydrodynamics in cylindrical coordinates. *ANZIAM*, 44:C114–C130, 2001.
114. L. D. Landau, L. P. Pitaevskii, A. M. Kosevich, and E. M. Lifshitz. *Theory of Elasticity*. Butterworth Heinemann, Jordan Hill, Oxford, 1986.
115. R. Stellingwerf and C. A. Wingate. Impact modeling with sph. *Los Alamos Nat. Lab. Rep.*, 1992.
116. J. J. Monaghan. Simulating free surface flows with sph. *J. Comput. Phys.*, 110:399–406, 1994.
117. G. Oger, M. Doring, B. Alessandrini, and P. Ferrant. Sph: Towards the simulation of wave-body interactions in extreme seas, 2004.
118. J. Kajtar and J.J. Monaghan. Sph simulations of swimming linked bodies. *J. Comput. Phys.*, 227:8568–8587, 2008.
119. J. J. Monaghan and A. Kos. Scott russell’s wave generator. *Phys. of Fluids*, 12:622–630, 2000.
120. M. Yildiz, R. A. Rook, and A. Suleman. Sph with the multiple boundary tangent method. *Int. J. Numer. Methods Engrg.*, (10).
121. J. P. Morris, P. Fox, and Y. Zhu. Modeling low reynolds number incompressible flows using sph. *J. Comput. Phys.*, 136:214–226, 1997.
122. R. I. Klein, R. Fisher, and C. F. McKee. Sph simulations of swimming linked bodies. *RevMexAA*, 22:3–7, 2004.
123. B. Van Leer. Towards the ultimate conservative difference scheme ii. monotonicity and conservation combined in a second order scheme. *J. Comput. Phys.*, 14:361–370, 1974.
124. G. A. Sod. A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws. *J. Comput. Phys.*, 27:1–31, 1978.

125. W. F. Noh. *J. Comput. Phys.*, 72:78, 1987.
126. W. A. Newcomb. Lagrangian and hamiltonian methods in magnetohydrodynamics. *Nucl. Fus. Supp.*, 2:451–463, 1962.
127. S. J. Cummins and M. Rudman. An sph projection method. *J. Comput. Phys.*, 152:584–607, 1999.
128. E. Ziegler, K. Dolag, and M. Bartelmann. Divergence cleaning techniques in smoothed particle magnetohydrodynamics. *Astron. Nachr.*, 327:607–608, 2006.
129. A. K. Chaniotis, D. Poulikakos, and P. Koumoutsakos. Remeshed smoothed particle hydrodynamics for the simulation of viscous and heat conducting flows. *J. Comput. Phys.*, 182:67–90, 2002.
130. J. Bonet and S. Kulasegaram. A simplified approach to enhance the performance of smooth particle hydrodynamics methods. *Appl. Math. and Computation*, 126:133–155, 2002.
131. H. F. Schwaiger. An implicit corrected sph formulation for thermal diffusion with linear free surface boundary conditions. *Int. J. Numer. Methods Engrg.*, 75:647–671, 2008.
132. P. P. Chinchapatnam, K. Djidjeli, and P. B. Nair. Unsymmetric and symmetric meshless schemes for the unsteady convection-diffusion equation. *Comput. Methods Appl. Mech. Engrg.*, 195:2432–2453, 2006.
133. S. N. Atluri and T. Zhu. A new meshless local petrov-galerkin (mlpg) approach in computational mechanics. *Comp. Mech.*, 22:117–127, 1998.
134. S. N. Atluri, Z. D. Han, and A. M. Rajendran. A new implementation of the meshless finite volume method, through the mlpg mixed approach. *Comput. and Struct.*, 6:491–513, 2004.
135. T. Zhu and S. N. Atluri. A modified collocation method and a penalty formulation for enforcing the essential boundary conditions in the element free galerkin method. *Comp. Mech.*, 21:211–222, 1998.
136. S. N. Atluri, H. T. Liu, and Z. D. Han. Meshless local petrov galerkin (mlpg) mixed collocation method for elasticity problems. *Comput. and Struct.*, 14:141–152, 2006.
137. A. L. Wilmot-Smith. Magnetic diffusion. 2003.
138. A. L. Wilmot-Smith and E. R. Priest. Magnetic diffusion and the motion of field lines. *Geophys. and Astrophys. Fl. Dyn.*, 99(2):177–197, 2005.
139. D. S. Balsara. Total variation diminishing scheme for adiabatic and isothermal magnetohydrodynamics. 116:133, 1998.
140. S. A. Orszag and C. M. Tang. Small-scale structure of two-dimensional magnetohydrodynamic turbulence. *J. Fl. Mech.*, 90:129–143, 1979.
141. K. L. Baker, R. D. Horton, D. Q. Hwang, R. W. Evans, and S. Howard. Simulation of the blowby effect in an accelerating compact toroid. *IEEE Trans. Plasma Sci.*, 30:48–49, 2002.
142. D. Q. Hwang, H. S. McLean, K. L. Baker, R. W. Evans, R. D. Horton S. D. Terry, S. Howard, and G. L. Schmidt. Interaction of a spheromak-like compact toroid with a high beta spherical tokamak plasma. *Nuclear Fusion*, 40:897–905, 2000.
143. Visit web page.
144. J. N. Johnson, D. Q. Hwang, R. D. Horton, R. W. Evans, and J. M. Owen. Simulated and experimental compression of a compact toroid (accepted). *Appl. Phys. Letters*, 95:11501, 2009.
145. R. W. Hockney and J. W. Eastwood. *Computer Simulation Using Particles*. Taylor and Francis, 1989.
146. J. Buckheit and D. L. Donoho. *Wavelab and reproducible research*. Springer Verlag, 1995.

Appendix A

Specific Forms of MLPG Matrices

In this section we give some specific forms of the MLPG matrices we discussed in Chapter 4. We obtain explicit expressions for (4.45) in terms of the test and trial functions and their partial derivatives for the MLPG1 variant.

Consider a trial function $\phi_i(\mathbf{x})$. Since electric and magnetic fields are always three-dimensional (even in problems with one-dimensional symmetry), we can treat the function as a scalar function multiplied by the 3×3 identity tensor. Then, if we approximate a vector field \mathbf{F} by

$$\mathbf{F}(\mathbf{x}) = \sum_j \phi_j \hat{\mathbf{F}}_j \quad (\text{A.1})$$

and its curl by

$$\nabla \times \mathbf{F}(\mathbf{x}) \equiv \sum_j \nabla \times \phi_j \cdot \hat{\mathbf{F}}_j \quad (\text{A.2})$$

we can equate the components of the curl of (A.1) to the components of (A.2):

$$[\nabla \times \mathbf{F}(\mathbf{x})]_x = \sum_{j \in \Omega_i} \phi_{j,y} \hat{F}_z - \phi_{j,z} \hat{F}_y \quad (\text{A.3})$$

$$[\nabla \times \mathbf{F}(\mathbf{x})]_y = \sum_{j \in \Omega_i} \phi_{j,z} \hat{F}_x - \phi_{j,x} \hat{F}_z \quad (\text{A.4})$$

$$[\nabla \times \mathbf{F}(\mathbf{x})]_z = \sum_{j \in \Omega_i} \phi_{j,x} \hat{F}_y - \phi_{j,y} \hat{F}_x \quad (\text{A.5})$$

and can infer that

$$\nabla \times \phi_j(\mathbf{x}) = \begin{pmatrix} 0 & -\phi_{j,z}(\mathbf{x}) & \phi_{j,y}(\mathbf{x}) \\ \phi_{j,z}(\mathbf{x}) & 0 & -\phi_{j,x}(\mathbf{x}) \\ -\phi_{j,y}(\mathbf{x}) & \phi_{j,x}(\mathbf{x}) & 0 \end{pmatrix} \quad (\text{A.6})$$

where the subscripts $_{,x}$, $_{,y}$, and $_{,z}$ refer to the x , y and z derivatives of the trial functions.

A similar analysis of $\mathbf{n} \times \mathbf{F}$ yields

$$\mathbf{n} \times \cdot = \begin{pmatrix} 0 & -n_z & n_y \\ n_z & 0 & -n_x \\ -n_y & n_x & 0 \end{pmatrix} \quad (\text{A.7})$$

To obtain an expression for $\mathbf{n} \times \nabla \times \phi_j$, we may take the matrix product of (A.7) with (A.6). We can use the same method to derive an expression for the curl of the test function $\psi_i(\mathbf{x})$.

Using these definitions of the curl for both the trial and test functions, we can rewrite the matrices and the load vector in (4.45) in terms of the shape functions and their partial derivatives, clarifying the analysis. The explicit 3×3 form of the mass matrix is

$$\mathbf{M}_{ij} = \int_{\Omega_i} \begin{pmatrix} \psi_i \phi_j & 0 & 0 \\ 0 & \psi_i \phi_j & 0 \\ 0 & 0 & \psi_i \phi_j \end{pmatrix} d\Omega$$

The stiffness matrix is

$$\begin{aligned} \mathbf{K}(\boldsymbol{\eta})_{ij} &= \int_{\Omega_i} \begin{pmatrix} 0 & -\psi_{i,z} & \psi_{i,y} \\ \psi_{i,z} & 0 & -\psi_{i,x} \\ -\psi_{i,y} & \psi_{i,x} & 0 \end{pmatrix} \cdot \begin{pmatrix} \eta_{xx} & 0 & 0 \\ 0 & \eta_{yy} & 0 \\ 0 & 0 & \eta_{zz} \end{pmatrix} \cdot \begin{pmatrix} 0 & -\phi_{j,z} & \phi_{j,y} \\ \phi_{j,z} & 0 & -\phi_{j,x} \\ -\phi_{j,y} & \phi_{j,x} & 0 \end{pmatrix} d\Omega \\ &- \int_{\Gamma_i \cup (\Gamma_B \cup \Gamma_K)} \psi_i \begin{pmatrix} 0 & -n_z & n_y \\ n_z & 0 & -n_x \\ -n_y & n_x & 0 \end{pmatrix} \cdot \begin{pmatrix} 0 & -\phi_{j,z} & \phi_{j,y} \\ \phi_{j,z} & 0 & -\phi_{j,x} \\ -\phi_{j,y} & \phi_{j,x} & 0 \end{pmatrix} d\Gamma \end{aligned} \quad (\text{A.8})$$

The load vector is

$$\mathbf{f}_i = - \int_{\Gamma_i \cup (\Gamma_B \cup \Gamma_K)} \begin{pmatrix} n_y E_{\Gamma_z} - n_z E_{\Gamma_y} \\ n_z E_{\Gamma_x} - n_x E_{\Gamma_z} \\ n_x E_{\Gamma_y} - n_y E_{\Gamma_x} \end{pmatrix} d\Gamma \quad (\text{A.9})$$

Appendix B

Software Environment

It has been said [146] that “an article about computational science in a scientific publication is not the scholarship itself, it is merely advertising of the scholarship.” The importance of reproducible research is beginning to hit the computational science community in a significant way. Indeed, the January/February 2009 issue of *Computing in Science and Engineering* is devoted almost entirely to this subject.

All of the test problems and simulations discussed in this thesis were run using *Spasmos*, a 1D/2D/3D resistive MHD code in which I have implemented the numerical methods discussed. This code is freely available to the interested reader, available for download from Sourceforge.net at <http://spasmos.sf.net> using Subversion. The version of *Spasmos* used to run these problems is tagged as *thesis*.

Spasmos is a set of Python libraries implemented in C and C++ and builds and runs on several modern platforms. It uses several freely-available software libraries and attempts to be reasonably simple to install and use. While no user’s manual exists at this time, the Python and C documentation is extensive, and the intrepid computational scientist should find several examples of problems in the source code to adapt to his or her needs. A seasoned Python programmer should find the idioms very familiar.

Simulation input is written as a Python script. Python is an interpreted language much like Matlab (see <http://www.python.doc> for an introduction). The scripts for the test problems in this thesis can be found in the `docs/thesis/testProblems` directory of the source tree. To run a script, simply pass it as a

parameter to the Python interpreter, e.g.

```
python testProblem.py
```

Each of these scripts can be invoked with a number of command-line parameters that dictate how its simulation should be run. The available parameters can be viewed by using the `--help` option for a particular script. For convenience, I list the scripts and the problems to which they correspond here.

integrateLinearDiffEq.py The convergence study for the various time integration methods discussed in Chapter 3.

hydrostatic-equilibrium.py The convergence study for hydrostatic equilibrium in one, two, and three spatial dimensions.

acoustic-wave-1d.py The convergence study for acoustic waves propagating in one dimension.

Sod.py The convergence study for the Sod shock tube problem.

Noh-planar.py The convergence study for the planar Noh problem in one dimension.

Noh-cylindrical.py The convergence study for the cylindrical Noh problem in two dimensions.

Noh-spherical.py The convergence study for the spherical Noh problem in three dimensions.

magnetic-decay-1d.py The convergence study for the decay of a magnetic pulse in one dimension.

magnetic-decay-2d.py The convergence study for the decay of a magnetic pulse in two dimensions.

inhomogeneous-mag-diff-1d.py Convergence studies for magnetic diffusion in conductors with spatially- and temporally-varying resistivities in one dimension.

magnetic-flux-tube-decay-2d.py The convergence study for the decay of a magnetic flux tube in two dimensions.

magnetic-advection-1d.py The convergence study for the advection of a magnetic pulse with and without magnetic diffusion.

MHD-wave-1d.py The convergence study for the propagation of magnetosonic waves in one dimension.

circularly-polarized-alfven-wave-1d.py The convergence study for the propagation of Alfvén waves in one dimension.

magnetic-shock-tube-1d.py Convergence studies for magnetized shock tube problems, including Brio and Wu's problem and the magnetic compression and rarefaction problems.

Balsara-blast-wave.py The convergence study for Balsara's spherical MHD blast wave problem in two dimensions.

Orszag-Tang-vortex-2d.py The convergence study for the Orszag-Tang vortex problem in two dimensions.